Operations and Process Management

Principles and practice for strategic impact

Fifth edition

Nigel Slack and
Alistair Brandon-Jones

Pearson

OPERATIONS AND PROCESS MANAGEMENT

OPERATIONS AND PROCESS MANAGEMENT

Principles and Practice for Strategic Impact

Fifth Edition

Nigel Slack

Alistair Brandon-Jones

Pearson

Harlow, England • London • New York • Boston • San Francisco • Toronto • Sydney • Dubai • Singapore • Hong Kong
Tokyo • Seoul • Taipei • New Delhi • Cape Town • São Paulo • Mexico City • Madrid • Amsterdam • Munich • Paris • Milan

PEARSON EDUCATION LIMITED
Kao Two
Kao Park
Harlow CM17 9NA
United Kingdom
Tel: +44 (0)1279 623623
Web: www.pearson.com/uk

First published 2006 (print)
Second edition published 2009 (print)
Third edition published 2012 (print and electronic)
Fourth edition published 2015 (print and electronic)
Fifth edition published 2018 (print and electronic)

ISBN: 978-1-292-17613-0 (print)
978-1-292-17617-8 (PDF)
978-1-292-17618-5 (ePub)

British Library Cataloguing-in-Publication Data
A catalogue record for the print edition is available from the British Library

Library of Congress Cataloging-in-Publication Data
Names: Slack, Nigel, author. | Brandon-Jones, Alistair, author.
Title: Operations and process management principles and practice for strategic impact / Nigel Slack, Alistair Brandon-Jones.
Description: Fifth edition. | United Kingdom : Pearson Education Limited, [2018] | Includes bibliographical references and index.
Identifiers: LCCN 2017050767| ISBN 9781292176130 (Print) | ISBN 9781292176178 (PDF) | ISBN 9781292176185 (ePub)
Subjects: LCSH: Manufacturing processes. | Process control.
Classification: LCC TS183 .O64 2018 | DDC 658.5–dc23
LC record available at https://lccn.loc.gov/2017050767

10 9 8 7 6 5 4 3 2 1
22 21 20 19 18

Front cover image © PM Images/Getty Images
Print edition typeset in 9.25/12.5pt Syntax Com by SPi Global
Printed in Slovakia by Neografia

NOTE THAT ANY PAGE CROSS REFERENCES REFER TO THE PRINT EDITION

Brief contents

Contents

Supporting resources

Visit www.pearsoned.co.uk/slack to find valuable online resources

For lecturers

- Instructor's Manual, including teaching notes for each chapter, model answers to all questions in the book and guided solutions for all case studies.
- PowerPoint slides

For students

- The student companion website provides model answers to the first two 'Applying the Principles' in each chapter.
  For more information, please contact your local Pearson Education sales representative or visit www.pearsoned.co.uk/slack.

Guide to case studies

Chapter	Case name and description	Region	Manufacturing/ service	Company size	Topics/techniques
Chapter 1 Operations and processes	Design house partnerships at concept design services	Europe	M,S	Medium	Role of operations, process objectives, types of operation and process
Chapter 2 Operations and strategic impact	McDonald's: half a century of growth	World	S	Large	Operations strategy, operations objectives, strategic fit
Chapter 3 Product and service innovation	Developing 'Savory Rosti-crisps' at Dreddo Dan's	World	M	Large	Product development, operations strategy, process performance
Chapter 4 Operations scope and structure	Aarens Electronic	Europe	S/M	Medium	Location, capacity, scope of operations
Chapter 5 Process design 1 – Positioning	McPherson Charles Solicitors	UK	S	Medium	Process design, job design, process technology, process resourcing
Chapter 6 Process design 2 – analysis	The Action Response Applications Processing Unit (ARAPU)	Africa, Asia, UK	S	Small	Process design, process mapping, balancing, Little's law
Chapter 7 Supply chain management	Supplying fast fashion	World	S M	Large	Outsourcing, supply chain design, fast response
Chapter 8 Capacity management	Blackberry Hill Farm	UK	S M	Small	Capacity management, forecasting, cumulative production and demand plotting
Chapter 9 Inventory management	supplies4medics.com	Europe	S	Medium	Inventory management, Inventory information systems, ABC analysis
Chapter 10 Resource planning and control	subText Studios Singapore	Singapore	S	Medium	Planning and control, Gantt charts, activity monitoring, controlling activities
Chapter 11 Lean synchronisation	Saint Bridget's Hospital	Europe	S	Medium	Improvement, quality, application of lean principles
Chapter 12 Improvement	Ferndale Sands Conference Centre	Australia	S	Small	Improvement, performance, prioritisation
Chapter 13 Quality management	Turnround at the Preston plant	Canada	M	Medium	Improvement principles, statistical process control, process learning, operations capabilities
Chapter 14 Risk and resilience	Slagelse Industrial Services (SIS)	Denmark	S M	Large	Risk, failure prevention, supplier selection, relationship management
Chapter 15 Project management	United Photonics Malaysia Sdn Bhd	Malaysia	S	Large	Project planning, project risk, project monitoring

Preface

Why is operations and process management essential?

Because it is about getting things done. Because without effective operations and processes there can be no long-term success for any organisation. Because it is at the heart of what all organisations do; they create value through their productive resources. Because it is the essential link that connects broad long-term strategy and day-to-day ongoing activities. This is why operations and process management has been changing. It has always been exciting, and it has always been challenging, but now it has acquired a much more prominent profile. The current edition reflects this in a number of ways.

It stresses the importance of operations and process management

Of course, it has always been important, but increasingly managers in all types of enterprise are accepting that operations management can make or break their businesses. Effective operations management can keep costs down, enhance the potential to improve revenue, promote an appropriate allocation of capital resources and, most important, develop the capabilities that provide future competitive advantage.

It stresses the real strategic impact of operations and process management

Operations are not always operational. The operations function also has a vital strategic dimension, and operations management is now expected to play a part in shaping strategic direction, not just responding to it.

It stresses that operations and process management matters to all sectors of the economy

At one time operations management was seen as being most relevant to manufacturing and a few types of mass service businesses. Now the lessons are seen applying to all types of enterprise; all types of service and manufacturing, large or small organisations, public or private, for-profit or not-for-profit.

It stresses that operations and process management is of interest to all managers

Perhaps most importantly, because operations management is accepted as being founded on the idea of managing process, and because managers in all functions of the business are now accepting that they spend much of their time managing processes, it is clear that to some extent, all managers are operations managers. The principles and practice of operations management are relevant to every manager.

It extends the scope of operations and process management

The obvious unit of analysis of operations management is the operations function itself – the collection of resources and processes that produce products and services. But, if managers from other functions are to be included, operations management must also address itself to process management at a more generic level. Also, no operation can consider itself in isolation from its customers, suppliers, collaborators and competitors. It must see itself as part of the extended supply network. Operations management increasingly needs to work at all three levels of analysis – the individual process, the operation itself and the supply network.

All this has implications for the way operations management is studied, especially at post-experience and postgraduate levels, and the way operations management is practised. It has also very much shaped the way this book has been structured. In addition to covering all the important topics that make the subject so powerful, it places particular emphasis on the following:

- **Principles** – that is, the core ideas that describe how operations behave, how they can be managed and how they can be improved. These are not immutable laws or prescriptions that dictate how operations *should* be managed, nor are they descriptions that simply explain or categorise issues. But they are indications of important underlying ideas.
- **Diagnosis** – an approach that questions and explores the fundamental drivers of operations performance. Aims to uncover or 'diagnose' the underlying trade-offs that operations need to overcome and the implications and consequences of the courses of action that could be taken.
- **Practice** – anyone with managerial experience, or who is approaching careers choices, understands the importance of developing practical knowledge and skills that can be applied in practice. This requires an approach, as well as frameworks and techniques, which can be adapted to take account of the complexity and ambiguity of operations, yet give guidance to identifying and implementing potential solutions.

Who should use this book?

This book is intended to provide an introduction to operations and process management for everyone who wishes to understand the nature, principles and practice of the subject. It is aimed primarily at those who have some management experience (although no prior academic knowledge of the area is assumed), or who are about to embark on a career in management. For example:

- *MBA students* should find that its practical discussions of operations management activities enhance their own experience.
- *Postgraduate students* on other specialist masters degrees should find that it provides them with a well-grounded and, at times, critical approach to the subject.
- *Executives* should find its diagnostic structure helpful to provide an understandable route through the subject.

What makes this book distinctive?

It has a clear structure

The book is structured on a model of operations management that distinguishes between activities that contribute to the direction, design, delivery and development of operations and processes.

It uses diagnostic logic chains

Every chapter follows a series of questions that forms a 'diagnostic logic' for the topic. These are the questions that anyone can ask to reveal the underlying state of their, or any other, operations. The questions provide an aid to diagnosing where and how an operation can be improved.

It is illustrations-based

Operations management is a practical subject and cannot be taught satisfactorily in a purely theoretical manner. Because of this, each chapter starts with two real-life examples of how the topic is treated in practice and provides additional examples in relation to specific issues within each chapter.

It identifies key operations principles

Whenever a core idea of operations and process management is described in the text, a brief 'operations principle' summary is included in the margin. This helps to distil those essential points of the topic.

It includes critical commentaries

Not everyone agrees about what is the best approach to the various topics and issues within the subject. This is why we have, at the end of each chapter, included a 'critical commentary'. These are alternative views to the one being expressed in the main flow of the text. They do not necessarily represent our view, but they are worth debating.

Each chapter includes summary checklists

Each chapter is summarised in the form of a list of checklist questions. These cover the essential questions that anyone should ask if they wish to understand the way their own or any other operation works. More importantly, they can also act as prompts for operations and process improvement.

Each chapter finishes with a case study

Every chapter includes a case study, relating real or realistic situations that require analysis, decision, or both. The cases have sufficient content to serve as the basis of case sessions in class, but are short enough to serve as illustrations for the less formal reader.

Each chapter includes an 'applying the principles' section

Selected problems, short exercises and activities are included at the end of each chapter. These provide an opportunity to test out your understanding of the principles covered in the chapter.

Each chapter includes a 'taking it further' section

A short annotated list of further reading and useful websites is provided which takes the topics in the chapter further, or treats some important related issues.

Suggested 'model answers' are available for all the 'applying the principles' exercises

Answers to the first two questions are available on the companion website for students. Answers to all the questions are available to bone fide tutors and lecturers.

Instructor's manual and PowerPoint slides

Visit www.pearsoned.co.uk/slack to find valuable online resources. A dedicated updated web-based instructor's manual is available to lecturers adopting this textbook. It includes teaching notes for all chapters, guided solutions for all case studies in the book, guided solutions for active cases and ideas for teaching them. A set of PowerPoint slides featuring figures and illustrations from the main text is also available.

About the authors

NIGEL SLACK is an Emeritus Professor of Operations Management and Strategy at Warwick University, an Honorary Professor at Bath University and an Associate Fellow of Said Business School, Oxford University. Previously he has been Professor of Service Engineering at Cambridge University, Professor of Manufacturing Strategy at Brunel University, a University Lecturer in Management Studies at Oxford University and Fellow in Operations Management at Templeton College, Oxford. He worked initially as an industrial apprentice in the hand-tool industry and then as a production engineer and production manager in light engineering. He holds a Bachelor's degree in Engineering and Master's and Doctor's degrees in Management, and is a Chartered Engineer. He is the author of many books in Operations Management, including The Operations Advantage, published in 2017, Operations Management (with Alistair Brandon-Jones and Robert Johnston), the eighth edition published in 2016, Operations and Process Management (with Alistair Brandon-Jones), the fifth edition published in 2018, Operations Strategy (with Michael Lewis), the fifth edition published in 2017, The Manufacturing Advantage, published in 1991, Making Management Decisions (with Steve Cooke), published in 1991, Service Superiority (with Robert Johnston), published in 1993, The Blackwell Encyclopedic Dictionary of Operations Management, published in 1997 and Perspectives in Operations Management (with Michael Lewis), published in 2003. Nigel has authored numerous academic papers and chapters in books. He also acts as a consultant to many international companies around the world in many sectors, especially financial services, transport, leisure and manufacturing. His research is in the operations and manufacturing flexibility and operations strategy areas.

ALISTAIR BRANDON-JONES is a Full Chaired Professor in Operations and Supply Management at the University of Bath, and an Adjunct Professor at Hult International Business School. He was formerly a Reader at Manchester Business School, an Assistant and Associate Professor at Bath University and a Teaching Fellow Warwick Business School, where he also completed his PhD. In addition to Essentials of Operations Management, his other books include Operations and Process Management (with Nigel Slack), the fifth edition published in 2018, Operations Management (with Nigel Slack and Robert Johnston), the eighth edition published in 2016 and Quantitative Analysis in Operations Management (with Nigel Slack), published in 2008. Alistair is an active empirical researcher focusing on e-enabled operations and supply management, professional services and healthcare operations. He has published this research extensively in world-leading journals including Journal of Operations Management, International Journal of Operations and Production Management, International Journal of Production Economics and International Journal of Production Research. He has also disseminated his research through various practitioner publications, conferences, workshops and white papers. Alistair has consulting and executive development experience with organizations around the world, in various sectors including petrochemicals, health, financial services, manufacturing, defence and government. In addition, he has won numerous prizes for teaching excellence and contributions to pedagogy, including from the Times Higher Education, Association of MBAs (AMBA), Production Operations Management Society (POMS), University of Bath, University of Manchester and University of Warwick.

Acknowledgements

In preparing this book, the authors, as usual, unashamedly exploited their friends and colleagues. In particular we had invaluable help from a great and distinguished reviewer team and from colleagues who have provided valuable feedback on various aspects of the project. For their help we are particularly grateful to:

Malcolm Afferson, Sheffield Hallam University, UK
Pär Åhlström, Stockholm School of Economics, Sweden
David Bamford, University of Huddersfield, UK
Vikram Bhakoo, University of Melbourne, Australia
Mattia Bianchi, Stockholm School of Economics, Sweden
Umit Bitici, Strathclyde University, UK
Briony Boydell, Portsmouth University, UK
Raffaella Cagliano, Politecnico di Milano, Italy
Dan Chicksand, Aston University, UK
Ben Clegg, Aston University, UK
Paul Coughlan, Trinity College Dublin, Ireland
Pamela Danese, University of Padova, Italy
Marc Day, Reading University, UK
Stephen Disney, Cardiff University, UK
Des Doran, Kingston University, UK
Joy Field, Boton College, USA
Paul Forrester, Keele University, UK
Gino Franco, Derby University, UK
Abby Ghobadian, Reading University, UK
Janet Godsell, Warwick University, UK
Roger Hall, Huddersfield University, UK
Ingjaldur Hannibalsson, University of Iceland, Iceland
Linda Hendry, Lancaster University, UK
Matthias Holweg, Oxford University, UK
Koos Krabbendam, University of Twente, The Netherlands
Bart MacCarthy, Nottingham University, UK
John Maguire, Sunderland University, UK
Harvey Maylor, Cranfield University, UK
Duncan McFarlane, Cambridge University, UK
Ian McCarthy, Simon Fraser University, Canada
Ronnie Mcmillan, University of Strathclyde, UK
Phil Morgan, Oxford Brookes University, UK
Andy Neely, Cambridge University, UK
Steve New, Oxford University, UK
Carie Queenan, University of South Carolina, USA
Peter Race, Reading University, UK
Pietro Romano, Udine University, Italy
Sarah Schiffling, Lincoln University, UK
Mike Shulver, Birmingham University, UK
Alison Smart, University of Glasgow, UK
Rui Sousa, Universidade Católica Portuguesa, Portugal
Nigel Spinks, Reading University, UK
Martin Spring, Lancaster University, UK
Jan Olhager, Lund University, Sweden
Zoe Radnor, Loughborough University, UK
Venu Venugopal, Nyenrode University, The Netherlands
Ann Vereecke, Vlerick Business School, Belgium
Rohit Verma, Cornell Univerity, USA
Jan de Vries, University of Groningen, The Netherlands
Graham Walder, Wolverhampton University, UK
Helen Walker, Cardiff University, UK
Maggie Zeng, Gloucestershire University, UK

At Warwick, thanks go to, Vikki Abusidualghoul, Haley Beer, Nicola Burgess, Mehmet Chakkol, Max Finne, Emily Jamieson, Mark Johnson, Pietro Micheli, Giovanni Radaelli, Ross Ritchie, Rhian Silvestro and Chris Voss.

At Bath, thanks go to, Maria Battarra, Emma Brandon-Jones, Jie Chen, Günes Erdogan, Emmanuel Fragniere, Vaggelis Giannikas, Andrew Graves, Yufei Huang, Jooyoung Jeon, Adam Joinson, Richard Kamm, Mike Lewis, Sheik Meeran, Ibrahim Muter, Fotios Petropoulos, Lukasz Piwek, Tony Roath, Jens Roehrich, Brian Squire, Kate Sugar, Christos Vasilakis, Xingjie Wei, Emma Williams and Baris Yalabik.

The prerequisite for any book of this type is that it serves a real market need. We were privileged to receive advice from the Pearson team; some of the most insightful people in business education publishing, in particular, Natalia Jaszczuk, Catherine Yates, Carole Drummond, Emma Marchant and Rufus Curnow. Finally, to our families, Angela and Kathy, and Emma and Noah, your ongoing support is invaluable.

Nigel Slack and Alistair Brandon-Jones

Publisher's acknowledgements

We are grateful to the following for permission to reproduce copyright material:

Figures

Figure 2.8 from *The Operations Advantage*, Kogan Page (Slack, N. 2017) reproduced with permission; Figures 3.2, 4.4, 11.7, 15.2 from *Operations Management*, 8e, Pearson Education Ltd, UK (Slack, N.) Copyright © 2016; Figure 4.3 from *Operations Management*, 8e, Pearson Education Ltd, UK (Slack, N.) Copyright © 2016; Figures 4.6, 4.8 from *Operations Management*, 8e, Pearson Education Ltd, UK (Slack, N.) Copyright © 2016; Figure 7.5 adapted from Fisher, M.C. (1997) 'What is the right supply chain for your product?', *Harvard Business Review*, March–April, 105–116; Figure 7.8 after The 50,000 mile journey of Wimbledon's tennis balls, *WBS News* 02/07/2014, Copyright © 2017 Warwick Business School, used with permssion; Figure 13.8 from EFQM. Reproduced with permission.

Tables

Table 11.1 adapted from 'Making continual improvement: a competitive strategy for service firms', in Bowen, D.E., Chase, R.B., Cummings, T.G. and Associates (eds) *Service Management Effectiveness*, Jossey-Bass (Fitzsimmons, J.A. 1990).

Photos

Interview 10 from Interview with Joanne Cheung, Steve Deeley and other staff at Godfrey Hall, BMW Dealership, Coventry, used with permission; Interview 11 from Interview with Edward Kay, Tom Dyson and Olly Willans of Torchbox, used with permission; Case Study 11 from Janina Aarts and Mattia Bianchi, Department of Management and Organization, Stockholm School of Economics, used with permission; Interview 13 from Interview with Michael Purtill, General Manager, Four Seasons Canary Wharf Hotel, used with permission.

1

Operations and processes

Introduction

Operations management is about how organisations produce goods and services. Every product you use and every service you experience comes to you courtesy of the operations managers who organised its production. Though not always called operations mangers, they are the people who design, run and improve the processes that produce services and products for their customers, so effectively they are operations managers. But operations ***and process*** management is even wider than this. Managers in other functions, such as marketing, sales and finance, also manage processes. They supply internal 'customers' with services such as marketing plans, sales forecasts, budgets, and so on. In fact, all parts of organisations are made up of processes. That is why operations and process management is of direct relevance to all managers, irrespective of what type of organisation they work for, or which function they work in. And that is what this book is about – the tasks, issues and decisions that are necessary to manage processes effectively, both within the operations function, and in other parts of the business where effective process management is equally important. This is an introductory chapter, so we will examine some of the basic principles of operations and process management. The model developed to explain the subject is shown in Figure 1.1.

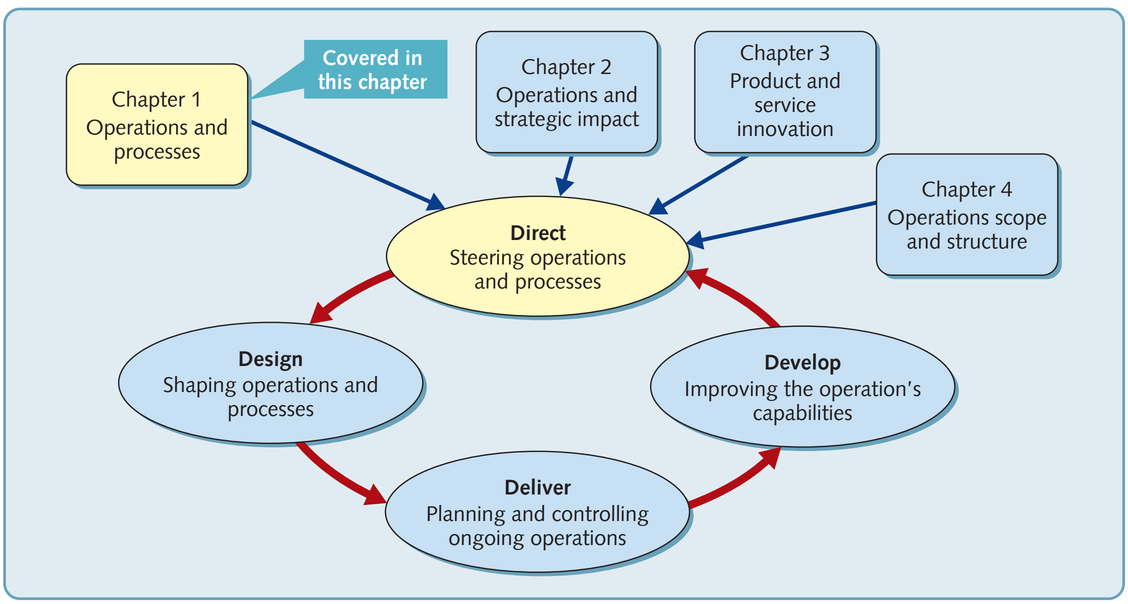

Figure 1.1 **Operations and process management is about how organisations produce goods and services**

EXECUTIVE SUMMARY

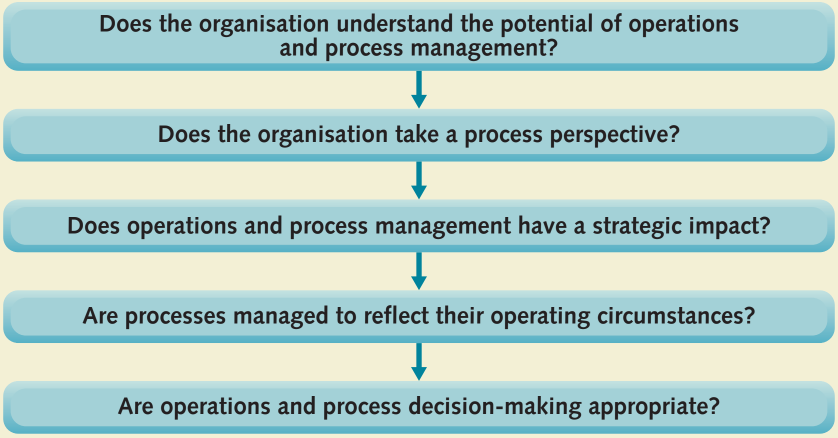

Each chapter has a short executive summary structured around the number of diagnostic questions used in the chapter. These diagnostic questions represent the basic line of enquiry that can reveal the nature and relevance of the topic covered in each chapter.

Does the organisation understand the potential of operations and process management?

The operations function is the part of the organisation that produces products or services. Every organisation has an operations function because every organisation produces some mixture of products and services. It is a central and important activity for *any* organisation. 'Operations' is not always called by that name, but whatever its name, it is always concerned with managing the core purpose of the business – producing some mix of products and services. Processes also produce products and services, but on a smaller scale. They are the component parts of operations. But other functions also have processes that need managing. In fact, *every* part of *any* business is concerned with managing processes. All managers have something to learn from studying operations and process management, because the subject encompasses the management of all types of operation, no matter in what sector or industry, and all processes, no matter in which function.

Does the business take a process perspective?

A 'process perspective' means understanding businesses in terms of all their individual processes. It is only one way of modelling organisations, but it is a particularly useful one. Operations and process management uses the process perspective to analyse businesses at three levels: the operations function of the business; the higher and strategic level of the supply network; and at a lower operational level of individual processes. Within the business, processes are only what they are defined as being. The boundaries of each process can be drawn as thought appropriate. Sometimes this involves radically reshaping the way processes are organised; for example, to form end-to-end processes that fulfil customer needs.

Does operations and process management have a strategic impact?

Operations and process management can make or break a business. Well-managed operations and processes can contribute to the strategic impact of the business in four ways: cost, revenue, investment and capabilities. Because the operations function has responsibility for much of a business's cost base, its first imperative is to keep costs under control. Additionally, it should look to enhance the business's ability to generate revenue through the way it provides service and quality. Furthermore, all failures are ultimately process failures; well-designed processes have less chance of failing and more chance of recovering quickly from failure. Because operations are often the source of much investment, it should aim to get the best possible return on that investment. Finally, the operations function should lay down the capabilities that will form the long-term basis for future competitiveness.

Should all processes be managed in the same way?

Not necessarily. Processes differ, particularly in what are known as the four Vs: volume, variety, variation and visibility. High-volume processes can exploit economies of scale and be systematised. High-variety processes require enough inbuilt flexibility to cope with the wide variety of activities expected of them. High-variation processes must be able to change their output levels to cope with highly variable and/or unpredictable levels of demand. High-visibility processes add value, while the customer is 'present' in some way and therefore must be able to manage customers' perceptions of their activities. Generally, high volume together with low variety, variation and visibility facilitate low-cost processes, while low volume together with high levels of variety, variation and visibility all increase process costs. Yet in spite of these differences, operations managers use a common set of decisions and activities to manage them. These activities can be clustered under four groupings: directing the overall strategy of the operation; designing the operation's products, services and processes; planning and controlling process delivery; and developing process performance.

Are operations and process decision-making appropriate?

The range of operations decisions are wide-ranging and cover four broad areas that we categorise as follows: '*directing* the overall strategy of the operation'; '*designing* the operation's processes'; 'planning and control process *delivery*'; and '*developing* process performance'. However, there are always overlaps and interrelationships between the categories. Yet, no matter what type of decision, operations managers use models (many of which are included in this book) to help them make decisions. Some models are quantitative, some are qualitative but, in practice, a blend of qualitative and quantitative approaches is often the most useful approach. Remember that all models are simplifications of a far more complex reality. This is one reason for the interest in 'behavioural operations management', which attempts to incorporate real (usually non-rational) behaviour into operations decision-making.

DIAGNOSTIC QUESTION

Does the organisation understand the potential of operations and process management?

Operations and process management is the activity of managing the resources and processes that produce products and services, for internal and external customers. It is a central and important activity for *any* organisation. The core body of knowledge for the subject comes from 'operations management', which examines how the 'operations function' of an organisation produces products and services for external customers. Some organisations may call an operations manager by some other name, for example, a 'fleet manager' in a logistics company, an 'administrative manager' in a hospital, or a 'store manager' in a supermarket. Note that we sometimes use the shorter terms 'the operation' or 'operations', interchangeably with the 'operations function'. Note also that we use 'the organisation', 'the business, or 'the enterprise' to describe any kind of human entity with a collective purpose.

All enterprises have 'operations'

All organisations have 'operations', because all organisations produce products, services, or some mixture of both. If you think that you do not have an operations function, you are wrong. If you think that your operations function is not important, you are also wrong. In most enterprises the operations function represents the bulk of its assets and the majority of its people; it is the means by which they serve their customers and provides an economic and/or social return for their stakeholders. An effective operations function has the potential to survive in a turbulent environment and the ability to maintain a steady improvement in its performance. By contrast, a poorly managed operations function, especially if it fails to provide an adequate service to its customers or fails to provide the efficiency to work within its cost constraints, will always prevent an organisation from achieving its objectives, whether social or economic.

OPERATIONS PRINCIPLE

All organisations have 'operations' that produce some mix of products and services.

But not all operations are the same

Look at the six businesses illustrated in Figure 1.2. There are two financial service companies, two manufacturing companies and two hotels. All of them have *operations functions* that produce the things that their customers are willing to pay for. Hotels produce accommodation services; financial services invest, store, move, or sell us money and investment opportunities; and manufacturing businesses physically change the shape and the nature of materials to produce products. These businesses are from different sectors (banking, hospitality and manufacturing), but it is not that they operate in different sectors of the economy that makes these businesses different from each other. There are often bigger differences *within* economic sectors than *between* them. The main difference between how their operations activities need to be managed is more closely related to the market position that they occupy. For example, all three operations in the left-hand column provide value-for-money products and services and compete largely on cost. The three in the right-hand column provide more 'up-market' products and services that are more expensive to produce and compete on some combination of high specification and customisation. The implication of this is important. It means

OPERATIONS PRINCIPLE

The economic sector of an operation is less important in determining how it should be managed than its intrinsic characteristics.

Figure 1.2 All types of business have 'operations' because all businesses produce some mix of products and services. The differences in the operations within a category of business are often greater than the differences between business sectors

that the surface appearance of a business and its economic sector are less important to the way its operations are managed than the intrinsic characteristics of what it is trying to achieve, including the volume of its output, the variety of products and services it produces and, above all, how it is trying to compete in its market.

Figure 1.3 illustrates how the scope of this subject has expanded over time. Originally, operations management was almost exclusively associated with the manufacturing sector – and usually called 'production' or 'manufacturing' management. Starting in the 1970s and 1980s, the term *operations management* reflected two trends. First, and most importantly, its use implied that many of the ideas, approaches and techniques traditionally used in the manufacturing sector could be equally applicable in the (much larger) service sector. The second use of the term expanded the scope of 'production' in manufacturing companies to include 'non-core', but important, production-related processes such as purchasing, physical distribution, after sales service, product development, and so on.

Operations *and process* management

More recently, use of the term *operations and process management* (or sometimes just process management) denotes the shift in the scope of the subject to include the whole organisation. Within any business, the production of products and services is not confined to the operations

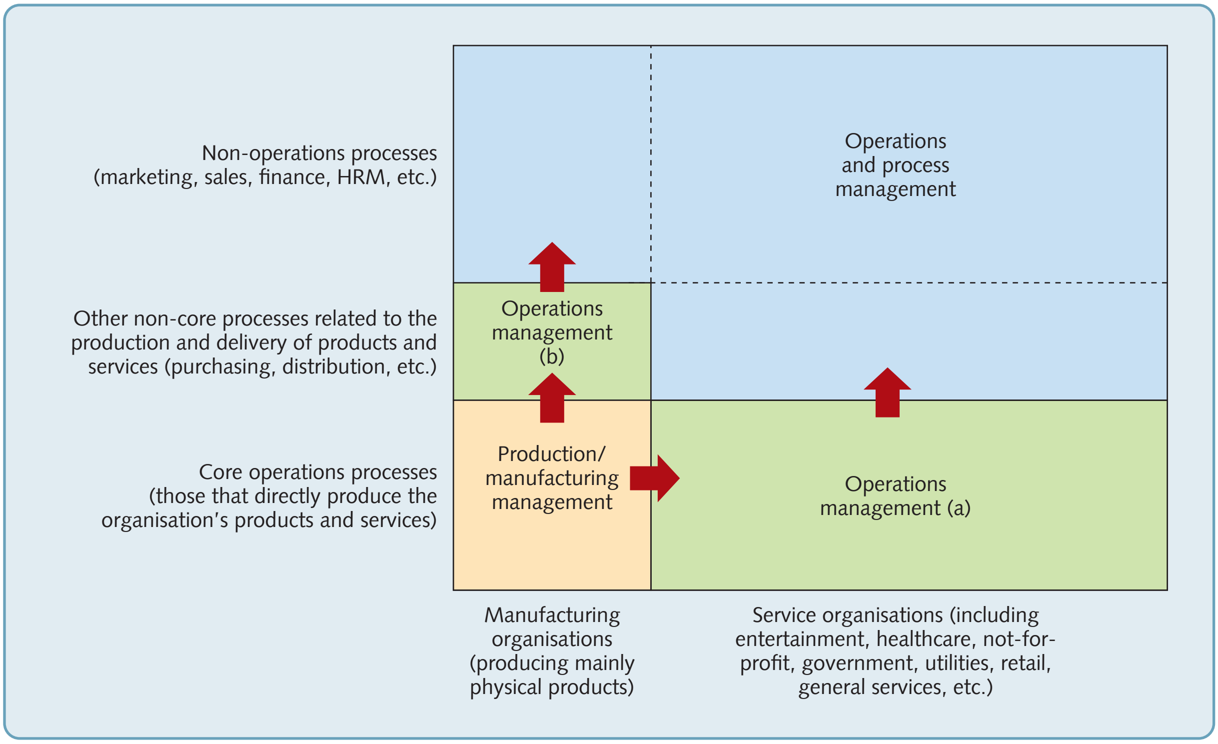

Figure 1.3 Operations management has expanded from treating only the core production processes in manufacturing organisations to include service organisations, non-core operations processes and processes in other functions such as marketing, finance and HRM

function. Every part of any business achieves their objectives by organising their resources such as people, information systems, buildings, and equipment into individual 'processes'. A 'process' is an arrangement of resources and activities that transform inputs into outputs that satisfy (internal or external) customer needs. For example, the marketing function 'produces' marketing plans and sales forecasts, the accounting function 'produces' budgets, the human resources function 'produces' development and recruitment plans, and so on. In fact, *every* part of *any* business is concerned with managing processes.

> **OPERATIONS PRINCIPLE**
> ***All operations are composed of processes. A process is an arrangement of resources and activities that transform inputs into outputs that satisfy (internal or external) customer needs.***

The difference between *operations* and *processes* is one of scale, and therefore complexity. Both transform inputs into outputs (we shall look at this idea later), but processes are the smaller version. They are the component parts – the building blocks – of an operation. So, 'operations and process management', is the term we use to encompass the management of all types of operation, no matter in what sector or industry, and all processes, no matter in which function of the business. The general truth is that processes are everywhere, and all managers have something to learn from studying operations and process management. This is very much how we treat the subject in this book and why it is called 'Operations and *Process* Management'. It includes an examination of the operations function in both manufacturing and service sectors, as well as the management of processes in non-operations functions.

At the beginning of each chapter, we present two examples of individual businesses, or types of business, that illustrate the topic examined in the chapter. Here we look at two businesses, one service company, and one manufacturing company, which have succeeded in large part because of their effective use of operations and process management principles.

EXAMPLE

Lego®[1]

Of all businesses, the toy business is one of the world's trickiest. It is difficult to forecast, unfailingly subject to fickle kids' latest fads and subject to constant technological innovation. Yet, in recent years, The LEGO Group, a privately held, family-owned company with headquarters in Billund, Denmark, has survived in the business, becoming one of the most reputable companies in the world, according to the Reputation Institute, and one of the leading manufacturers of play materials. It is a success founded on a deceptively simple idea. One Lego brick is unremarkable, but put one or two together and possibilities start to emerge. With another few bricks, the number of things you can create rises exponentially. There are more than 915 million possible ways of arranging six standard four-by-two bricks; and, however many bricks you assemble, irrespective of what colour or set they are from, your pieces will always fit together perfectly. All of the basic Lego elements use the same method to stick together. They have studs on top and tubes on the inside. The bricks studs are slightly bigger than the space between the tubes and the walls. Pressing the bricks together produces an 'interference fit' that provides a temporary joint without the use of an additional fastener. However, this principle does depend on very high levels of precision and quality in the manufacture of the elements, which explains the company's motto, 'Only the best is good enough'.

Lego bricks are manufactured at the Group's factories in Denmark, Hungary, The Czech Republic and Mexico, with the location of new factories built in Nyiregyhaza in Hungary and Jiaxing in China chosen to be near their key markets. Products made in these factories serve a global market. The aim, according to Bali Padda, Executive Vice President and Chief Operations Officer of The LEGO Group, is to 'build a stable manufacturing base around the world, ultimately making sure that LEGO products are available to children and their parents when and where they want them'. The company's operations processes are central to maintaining its reputation for quality and its ability to produce millions of elements profitably and sustainably. These processes start at the main warehouse, which contains the silos that hold raw plastic granulates. A complex arrangement of tubes links the silos to the moulding machines. The moulding stage is particularly important because of the need to make every Lego piece to a demanding level of precision, with tolerances as small as 10 micrometres. At each machine, the plastic is heated and pumped into the mould through a main channel, which divides into a number of narrower channels, each corresponding to a single brick and, when the plastic has solidified (after only a couple of seconds) the machines release the bricks into containers. A sensor detects when a container is full and a robot trolley automatically picks up boxes and leaves empty ones. The robots transport the boxes to conveyors, which move them into the storage area where robotic cranes stack them until needed. The automation means that few people are required for the process.

From there, some pieces go to the 'decoration' stage where they are individually painted. Decoration is the most expensive part of the Lego process. Other pieces go straight to packing, where the Lego sets take their final form. In the packaging process, the pieces go into a machine that separates them individually, counts them using optical sensors and places them in their box. The automatic movement system knows exactly how much a box should weigh at any stage and high-precision scales monitor the weight of the box as the packing process continues. Any deviation, even of a few micrograms, sets off an alarm. At the end of the process, the boxes are sealed shut, automatically weighed to ensure no components are missing, checked by a worker trained to look for errors (e.g. plastic bags sticking out of the box), packed six to a case by a robot and finally sent off for distribution.

Quality assurance staff perform frequent inspections and tests on the various Lego elements, such as drop, torque, tension, compression, bite and impact tests to make sure the toys are robust and safe. Only about eighteen of every million Lego elements produced (0.00002 per cent) fail to pass the tests. In addition, throughout the process, the company tries to achieve high levels of environmental sustainability. Plastic is extensively recycled in the factory. All scrap such as, for example, the plastic that fills the channels that take the hot plastic into moulds and faulty pieces that escape from automated handling, are ground up, and taken back into the production process. Similarly, the transparent plastic used to clean the channels when the production colour is changed in a moulding machine is also ground up and sold to companies that produce other plastic products.

EXAMPLE

Torchbox: award-winning web designers[2]

We may take it for granted, yet browsing websites, as part of your studies, your job, or your leisure, is an activity that we all do; probably every day, probably many times each day. It's important. All organisations need to have a web presence if they want to sell products and services, interact with their customers, or promote their cause. Not surprisingly, there is a whole industry devoted to designing websites so that they have the right type of impact. In fact, taken over the years, web development has been one of the fastest growing industries in the world. But it's also a tough industry. Not every web design company thrives, or even survives beyond a couple of years. To succeed, web designers need technology skills, design capabilities, business awareness and operational professionalism. One that has succeeded is Torchbox, an independently owned web design and development company based in Oxfordshire. Founded in 2000, it now employs 60 people, providing 'high-quality, cost-effective, and ethical solutions for clients who come primarily, but not exclusively, from the charity, non-governmental organisations and public sectors'.

Co-founder and Technical Director Tom Dyson has been responsible for the technical direction of all major developments. 'There are a number of advantages about being a relatively small operation', he said. 'We can be hugely flexible and agile, in what is still a dynamic market. But at the same time we have the resources and skills to provide a creative and professional service. Any senior manager in a firm of our size cannot afford to be too specialised. All of us here have their own specific responsibilities; however, every one of us shares the overall responsibility for the firm's general development. We can also be clear and focused on what type of work we want to do. Our ethos is important to us. We set out to work with clients who share our commitment to environmental sustainability and responsible, ethical business practice; we take our work, and that of our clients, seriously. If you're an arms dealer, you can safely assume that we're not going to be interested.'

Nevertheless, straightforward operational effectiveness is also essential to Torchbox's business. "*We know how to make sure that our projects run not only on time and to budget*" said Olly Willans, also a co-founder and the firms Creative Director, *'but we also like to think that we provide an enjoyable and stimulating experience – both for our customers' development teams and for our staff too. High standards of product and service are important to us: our clients want accessibility, usability, performance and security embedded in their web designs and, of course, they want things delivered on-time and on-budget. We are in a creative industry that depends on fast-moving technologies, but that doesn't mean that we can't also be efficient. We back everything we do with a robust feature-driven development process using a kanban project management methodology which helps us manage our obligations to our clients.'*

The 'kanban' approach used by the Torchbox web development teams originated from car manufacturers like Toyota (this is fully explained in Chapter 11). *'Using sound operations management techniques helps us constantly to deliver value to our clients*', said Tom Dyson, *'we like to think that our measured and controlled approach to handling and controlling work helps ensure that every hour we work produces an hour's worth of value for our clients and for us.'*

What do they have in common?

These two operations are very different. One is one of the best-known toy manufacturing companies in the world, the other is a small (but successful) company working in an industry that did not exist when Lego was founded. Yet the operations of both companies share the same basic objective

– managing the processes that produce their products and services in a manner that satisfies their customers, while making enough profit to thrive in the long term. Most of the managers in each company, irrespective of their formal title, will be concerned also with managing the processes that contribute to the success of their respective businesses. Of course, there are differences between each company's operations and processes, such as the scale at which they operate, the type of products and services they provide, the resources they use, and so on. But the managers in each company will be making the same *type* of decisions, even if *what* they actually decide is different. The fact that both companies are successful because of their innovative and effective operations and processes also implies further commonality. First, it means that they both understand the importance of taking a 'process perspective' in understanding their supply networks, running their operations, and managing all their individual processes. Without this, they could not have sustained their strategic impact in the face of stiff market competition. Second, both businesses will expect their operations to contribute to their overall competitive strategy. Third, in achieving a strategic impact, they both will have come to understand the importance of managing *all* their individual processes throughout the business so that they too can all contribute to the businesses, success.

DIAGNOSTIC QUESTION

Does the business take a process perspective?

Central to making operations and process management a significant contributor to a business's success is the idea of a 'process perspective'. If a business takes a process perspective, it understands that all parts of the business can be seen as processes, and that all processes can be managed using operations management principles. Yet, although important, a process perspective is not the only way of describing businesses, or any type of organisation. One could represent an organisation as a conventional 'organisational structure' that shows the reporting relationships between various departments or groups of resources. However, even a little experience in any organisation shows that rarely, if ever, does this fully represent the way the organisation actually works. Alternatively, one could describe an organisation through the way it makes decisions: how it balances conflicting criteria, weighs up risks, decides on actions and learns from its mistakes. Or, one could describe an organisation by explaining its culture – its shared values, ideology, pattern of thinking and day-to-day rituals, or its power relationships – how it is governed, seeks consensus (or at least reconciliation), and so on. Or, and this is the significant point, one can represent the organisation as a collection of processes, interconnecting and (hopefully) all contributing to fulfilling its strategic aims. This is the perspective that we emphasise throughout this book. As we define it here, the process perspective analyses businesses as a collection of interrelated processes. Some of these processes will be within the operations function, and will contribute directly to the production of its products and services. Other processes will be in the other functions of the business, but will still need managing using similar principles to those within the operations function.

None of these various perspectives on organisations gives a total picture. Each perspective adds something to our ability to understand, and therefore more effectively manage a business. Nor are these perspectives mutually exclusive. A process perspective does not preclude understanding the influence of power relationships on how processes work, and so on. We use the process perspective here, not because it is the *only* useful and informative way of understanding businesses, but because it is the perspective that directly links the way we manage resources in a business with its strategic impact. Without effective process management, the

OPERATIONS PRINCIPLE

There are many valid approaches to describing organisations. The process perspective is a particularly valuable one.

best strategic plan can never become reality. The most appealing promises made to clients or customers will never be fulfilled. In addition, the process perspective has traditionally been undervalued. The subject of operations and process management has only recently come to be seen as universally applicable and, more importantly, universally valuable.

So, operations and process management is relevant to all parts of the business

If processes exist everywhere in the organisation, operations and process management will be a common responsibility of all managers irrespective of which function they are in. Each function will have its 'technical' knowledge of course. In Marketing, this includes the market expertise needed for designing and shaping marketing plans; in Finance, it includes the technical knowledge of financial reporting conventions. Yet each will also have an *operations* role that entails using its processes to produce plans, policies, reports and services. For example, the marketing function has processes with inputs of market information, staff, computers, and so on. Its staff transform the information into outputs such as marketing plans, advertising campaigns and sales force organisations. In this sense, all functions are operations with their own collection of processes. The implications of this are very important. Because every manager in all parts of an organisation is, to some extent, an operations manager, they all should want to give good service to their (internal) customers, and they all should want to do this efficiently. Therefore, operations management must be relevant for all functions, units and groups within the organisation. And the concepts, approaches and techniques of operations management can be used to help improve any process in any part of the organisation.

OPERATIONS PRINCIPLE

All parts of the business manage processes so all parts of the business have an operations role and need to understand, operations management.

The 'input–transformation–output' model

OPERATIONS PRINCIPLE

All processes have inputs of transforming and transformed resources that they use to create products and services.

Central to understanding the processes perspective is the idea that all processes and operations transform *inputs* into *outputs*. Figure 1.4 shows the *general transformation process model* that is used to describe the nature of processes and operations. Put simply, processes and operations take in a set of input resources, some of which are transformed into outputs of products and/or services and some of which do the transforming.

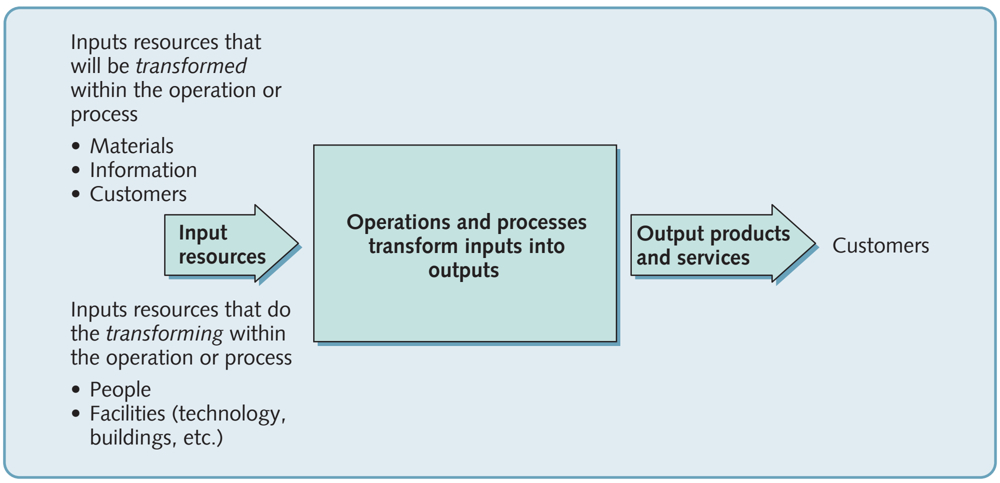

Figure 1.4 **All operations and processes are input–transformation–output systems that use 'transforming' resources to work on 'transformed' resources in order to produce products and services**

Input resources

Transformed resource inputs are the resources that are changed in some way within a process. They are usually materials, information or customers. For example, one process in a bank prints statements of accounts for its customers. In doing so, it is processing materials. In the bank's branches, customers are processed by giving them advice regarding their financial affairs, making payments, and so on . However, behind the scenes, most of the bank's processes are concerned with processing information about its customers' financial affairs. In fact, for the bank's operations function as a whole, its information-transforming processes are probably the most important. As customers, we may be unhappy with a clunky website or if we are not treated appropriately in the bank branch. But if the bank makes errors in our financial transactions, we suffer in a far more fundamental way.

> **OPERATIONS PRINCIPLE**
>
> ***Transformed resource inputs to a process are materials, information, or customers.***

> **OPERATIONS PRINCIPLE**
>
> ***All processes have transforming resources of facilities (equipment, technology, etc.) and people.***

Two types of *transforming* resource form the 'building blocks' of all processes. They are *facilities* – the buildings, equipment, plant and process technology of the operation; and *people* – who operate, maintain, plan and manage the operation.

The exact nature of both facilities and people will differ between processes. In a five-star hotel, facilities consist mainly of buildings, furniture and fittings. In a nuclear-powered aircraft carrier, its facilities are the nuclear generator, turbines and sophisticated electronic equipment. Although one operation is relatively 'low-technology' and the other 'high-technology', their processes all require effective, well-maintained facilities. Staff will also differ between processes. Most staff employed in a food processing plant may not need a very high level of technical skill, whereas most staff employed by an accounting firm in an audit process are highly skilled in their own particular 'technical' skill (accounting). Yet although skills vary, all staff have a contribution to make to the effectiveness of their operation. A food factory worker who does not follow hygiene procedures will damage the business and its customers just as surely as an accountant who cannot add up.

Output for customers

All processes and operations produce products and/or services, and although products and services are different, the distinction can be subtle. We will look at the differences between products and services later in this chapter. But whether classed as a product or service, the outputs from processes and operations are (or should be) intended to serve the requirements of customers. For the operations function of a business, the customers who need to be served are likely to be *external* customers – those parties or organisations that use the business's products or services. For the individual processes within a business, customers are usually (although not always) *internal* customers. Internal customers can be other processes, individuals, or groups within the business; in fact, anyone who is affected by the product or service.

The obvious difference between serving external and internal customers is their freedom of choice. External customers provide an organisation directly with its reason for being. They provide the revenue (directly or indirectly) on which the organisation depends. Consistently failing to satisfy external customers is likely to result in them going elsewhere. There is an obvious and clear reason for any operation to pay serious attention to how it serves external customers. Internal customers, at least in the short term, do not usually have the option of 'going elsewhere'. As the customer of your organisation's payroll process, you cannot choose to be served by an alternative payroll process if they make mistakes on your salary payments. Yet, although it may seem to have less impact initially, serving internal customers is, in the longer term, just

S suppliers	I inputs	P process	O output	C customers
Functional head Project manager Recruitment agency	Recruitment request Long-term resourcing plan Labour market reports	Agree essential and desirable characteristics ↓ Authorise budget and timescale ↓ Formally contract recruitment agency ↓ Agree shortlist of candidates ↓ Make offer and negotiate conditions ↓ Confirm start date and arrange induction	Qualified and experienced new staff member Completed staff records Payment authorisation	Functional head Project manager New recruit

Figure 1.5 A simple SIPOC analysis for a recruitment process performed by the Human Resources function of a company

as important. Poor service to internal customers will eventually result in poor service to external customers, or excessive costs to the organisation, or both.

SIPOC analysis

Do not dismiss the idea of the 'input–transformation–output' model as either too obvious or trivial, it can be the basis of a useful first step in understanding and improving processes. This is sometimes called SIPOC analysis. SIPOC stands for suppliers, inputs, process, outputs and customers. As its name implies, it is a method of formalising what a process requires as its inputs, where those inputs come from, what the process entails (at a relatively general rather than a detailed level), what the process produces as outputs and for whom the outputs are intended. Figure 1.5 shows a simple example that describes a recruitment process performed by the human resources function of a company. The advantage of such an analysis is that it helps all those involved in the process to understand (and, more important, agree) what it involves and where it fits within the business. More than this, it can prompt important questions that are sometimes overlooked. For example, exactly what information should suppliers to the process provide? What notice should they be required to give? In what form should they give the information? What are the important steps in the process and who is responsible for them? Who needs to be informed when the process is complete? How should its success be judged? And so on.

The three levels of analysis

Operations and process management uses the process perspective to analyse businesses at a number of levels – three in particular. The most obvious level is that of the business itself, or more specifically, the operations function of the business. (The other functions of the business could also be treated at this level.) And, while analysing the business at the level of 'the operation' is important, for a more comprehensive assessment we also need to analyse the contribution of operations and process management at a higher and more strategic level (the level of its supply network) and at a lower more operational level (the level of individual processes). These three levels of operations analysis are shown in Figure 1.6.

> OPERATIONS PRINCIPLE
>
> *A process perspective can be used at three levels; the level of the operation itself, the level of the supply network and the level of individual processes.*

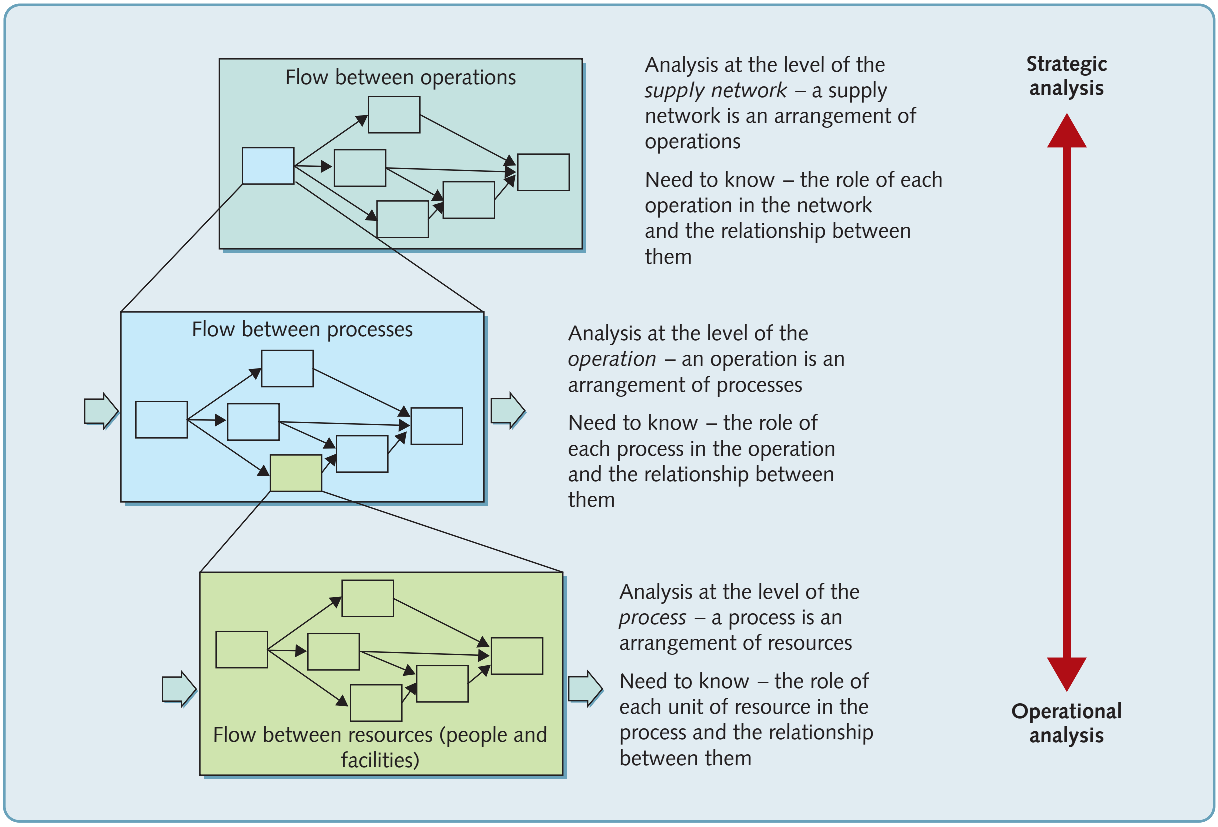

Figure 1.6 Operations and process management requires analysis at three levels: the supply network, the operation and the process

The process perspective at the level of the operation

The operations part of a business is itself an input–transformation–output system, which transforms various inputs to produce (usually) a range of different products and services. Table 1.1 shows some operations described in terms of their main resource inputs, the processes within the operations, and the outputs from the operation. Note how some of the inputs to the operation are transformed in some way, while other inputs do the transforming. Note also how in some operations customers themselves are inputs. (The airline, department store and police department are all like this.) This illustrates an important distinction between operations whose customers receive their outputs without seeing inside the operation, and those whose customers are inputs to the operation and therefore have some visibility of the operation's processes. Managing high-visibility operations where the customer is inside the operation usually involves different set of requirements and skills to those whose customers never see inside the operation. (We will discuss this issue of visibility later in this chapter.)

Most operations produce both products and services

The difference between a 'product' and a 'service' is not always obvious and this has provoked a lot of (not always useful) academic debate. Generally, a product is a physical and tangible thing (you can touch a television or a phone). A service is an activity that usually involves interaction with a customer (as with a physician) or something representing the customer (as with a package delivery service). Products can usually be stored for a time, some food products only for a few days and some buildings for thousands of years. The life of a service is often much

Table 1.1 Some operations described in terms of their inputs, purpose and outputs

Type of operation	*What are the operation's inputs?*	*What the processes within the operation?*	*What are operation's outputs?*
Airline	Aircraft Pilots and air crew Ground crew *Passengers* *Cargo*	Passenger check-in assistance Baggage drop Security / seat check Board passengers Fly passengers and freight around the world Flight scheduling In-flight passenger care Transfer assistance Baggage reclaim	Transported passengers and freight
Department store	*Goods for sale* Staff sales Computerised registers *Customers*	Source merchandise Manage inventory Display products Give sales advice Sales Aftercare Complaint handling Delivery service	Customers and goods 'assembled' together
Police department	Police officers Computer systems *Information* *Public (law-abiding and criminal)*	Crime prevention Crime detection Information gathering / collating Victim support Formally charging / detaining suspects, managing custody suites, liaising with court / justice system	Lawful society Public with feeling of security
Frozen food manufacturer	*Fresh food* Operators Food-processing equipment Freezers	Source raw materials Input quality checks Prepare ingredients Assemble products Pack products Fast freeze products Quality checks Finished goods inventory	Frozen food

Note: input resources that are transformed are printed in *italics*.

shorter. For example, the service of 'accommodation in a hotel room for tonight' will 'perish' if it is not sold before tonight – accommodation in the same room tomorrow is a different service.

Some operations produce just products and others just services, but most produce a mixture of the two. Figure 1.7 shows a number of operations positioned in a spectrum from almost 'pure' goods producers to almost 'pure' service producers. Crude oil producers are concerned almost exclusively with the product that comes from their oil wells. So too are aluminium smelters, but they might also produce some services such as technical advice. Machine tool manufacturers produce services such as technical advice and applications engineering services as well as products. The services produced by restaurants are an essential part of what the customer is paying for. They both manufacture food and provide service. An IT systems services company may produce software 'products', but more so, it is providing an advice and customisation service to its customers. A management consultancy, although producing reports and documents, would see itself largely as a service provider. Finally, some pure services do not produce products at all. A psychotherapy clinic, for example, provides therapeutic treatment for its customers without any physical product.

OPERATIONS PRINCIPLE

Most operations produce a mixture of tangible products and intangible services.

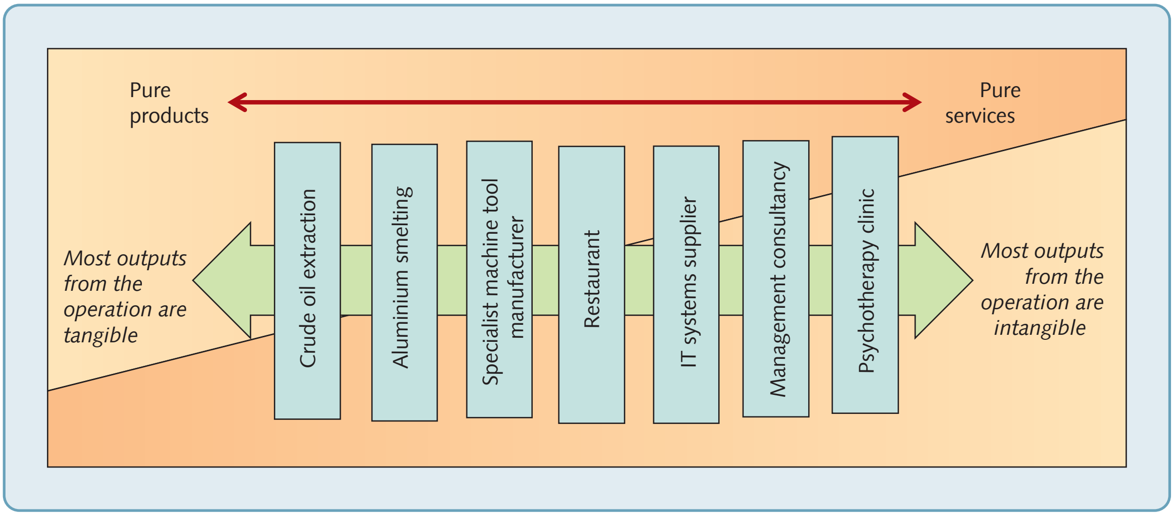

Figure 1.7 Relatively few operations produce either purely products or purely services. Most types of operation are a mix of goods and services

Services and products are merging

Increasingly, the distinction between services and products is seen as not particularly useful. Some authorities see the essential purpose of all businesses, and therefore all operations, as being to 'serve customers'. Therefore, they argue, all operations are service providers who may (or may not) produce physical products as a means of serving their customers. This idea, that all operations should be seen as offering 'value propositions' through service, has been called 'service-dominant logic'.[3] Amongst other things, it holds that service is the fundamental basis of exchange, that physical goods are simply the distribution mechanisms for the provision of service, and that the customer is always the co-creator of value. Our approach in this book is close to this in that we treat operations and process management as being important for all organisations. Whether they see themselves as manufacturers or service providers is very much a secondary issue.

The process perspective at the level of the supply network

Any individual operation is part of a greater network of operations. It will have operations that supply it with the products and services it needs to create its own products and services. Unless it deals directly with the end consumer, it will supply its customer operations who themselves will go on to supply their own customers. Moreover, any operation will probably have several suppliers, several customers and be in competition with other operations, with which it shares some suppliers and some customers. This collection of interconnected operations is called the supply network.

There are three important issues to understand about supply networks. First, they can be complex. Operations may have a large number of customers and suppliers who themselves have large numbers of customers and suppliers. The relationships between operations in the supply network can be subtle. One operation may be in direct competition with another in some markets, while at the same time acting as collaborators or suppliers to each other in others. Second, theoretically the boundaries of any operation's supply network can be very wide indeed. They could go back to the operation that digs raw material out of the ground and go forward to the ultimate reuse and/or disposal of a product. Sometimes it is necessary to consider these ultimate boundaries (for example, when studying the environmental sustainability of products), but generally some kind of boundary to the network needs to be set so that more attention can be given to the more immediate operations in the network. Third, supply networks are always changing. Not only do operations sometimes lose customers and win others, or change their suppliers, they also may acquire operations that once were their customers or suppliers, or sell parts of their business, so converting them into customers or suppliers.

Thinking about operations management in a supply network context is important. The overarching question for any operations manager is, 'Does my operation make a contribution to the supply network as a whole?' In other words, 'are we a good customer to our suppliers in the sense that the long-term cost of supply to us is reduced because we are easy to do business with?' And, 'are we good suppliers to our customers in the sense that, because of our understanding of the supply network as a whole, we understand their needs and have developed the capability to satisfy them?' Because of the significance of the supply network perspective we deal with it twice more in this book; at a strategic level in Chapter 4 where we discuss the overall structure and scope of individual operations, and at a more operational level in Chapter 7 where we examine the role of supply chains in the delivery of products and services.

The process perspective at the level of the individual process

Because processes are smaller versions of operations, they have customers and suppliers in the same way as whole operations. We can view any operation as a network of individual processes that interact with each other, with each process being, at the same time, an internal supplier and an internal customer for other processes. This 'internal customer' concept provides a model to analyse the internal activities of an operation. If the whole operation is not working as it should, we may be able to trace the problem back along this internal network of customers and suppliers. It can also be a useful reminder to all parts of the operation that, by treating their internal customers with the same degree of care that they exercise on their external customers, the effectiveness of the whole operation can be improved. Again, remember that many of an organisation's processes are not operations processes, but are part of some other function. Table 1.2 illustrates just some of the processes that are contained within some of the more common non-operations functions, the outputs from these processes and their 'customers'.

There is an important implication of visualising each function of an organisation as being a network of processes. The diverse parts of a business are connected by the relationships between their various processes; and the organisational boundaries between each function

Table 1.2 Some examples of processes in non-operations functions

Organisational function	*Some of its processes*	*Some outputs from its process*	*Customer(s) for its outputs*
Marketing and sales	Planning process Forecasting process Order-taking process	Marketing plans Sales forecasts Confirmed orders	Senior management Sales staff, planners, operations Operations, finance
Finance and accounting	Budgeting process Capital approval processes Invoicing processes	Budget Capital request evaluations Invoices	Everyone Senior management, requestees External customers
Human resources management	Payroll processes Recruitment processes Training processes	Salary statements New hires Trained employees	Employees All other processes All other processes
Information technology	Systems review process Help desk process System implementation project processes	System evaluation Advice Implemented working systems and aftercare	All other processes All other processes All other processes

and part of the business is really a secondary issue. Firms continually reorganise the boundaries between processes. They frequently move responsibility for tasks between departments. The tasks and the processes that perform them change less often. Similarly, tasks and processes may move between various businesses; that is what outsourcing and the 'do or buy' decision is all about (see Chapter 4). In other words, not only can separate businesses be seen as networks of processes, whole supply networks can also. Who owns which processes and how the organisational boundaries between them move are separate decisions.

OPERATIONS PRINCIPLE
Whole businesses, and even whole supply networks, can be viewed as networks of processes.

The 'line of sight' within process networks

Thinking about operations as networks of processes also gives a further benefit. It prompts the question of whether the people who operate each process have (what is known as) a clear 'line of sight' forward through to the external customers who will eventually be affected by their performance. If so, they will have a better chance of seeing how they contribute to the final value added for the operation's customers. Just as important, they will be better placed to help those other processes that lie between them and the customer to add value. Similarly, a clear 'line of sight' backwards helps to understand what is required from internal (and eventually external) suppliers. Certainly, a failure to understand how internal process chains work will reduce the effectiveness of the whole operation.

'End-to-end' business processes

It is worth remembering that what we choose to define as a specific process is not pre-determined. We can define what is inside a process in any way we want. The boundaries between processes, the activities they perform and the resources they use, are all there because they have been designed in that way. It is common in an organisation to find processes defined by the type of activity they engage in. For example, invoicing processes, product design processes, sales processes, warehousing processes, assembly processes, painting processes, and so on. This can be convenient because it groups similar resources together. But it is only one way of drawing the boundaries between processes. Theoretically, in large organisations there must be almost an infinite number of ways to collect activities and resources together as distinct processes. One way of redefining the boundaries and responsibilities of processes is to consider the 'end-to-end' set of activities that satisfy defined customer needs. Think about the various ways in which a business satisfies its customers. Many different activities and resources will probably contribute to 'producing' each of its products and services. Some authorities recommend grouping the activities and resources together in an end-to-end manner to satisfy each defined customer need. This approach is closely identified with the 'business process engineering' (or re-engineering) movement (examined in Chapter 12). It calls for a radical rethink of process design that will probably involve taking activities and resources out of different functions and placing them together to meet customer needs. Remember though, designing processes around end-to-end customer needs is only one way (although often the sensible one) of designing processes.

OPERATIONS PRINCIPLE
Processes are defined by how the organisation chooses to draw process boundaries.

EXAMPLE

The programme and video division (PVD)

A broadcasting company has several divisions including several television and radio channels (entertainment and news), a 'general services' division that includes a specialist design workshop, and the 'programme and video division' (PVD) that makes programmes and videos for a number of clients including the television and radio channels that are part of the same company. The original ideas for these programmes and videos usually come from the clients who commission them, although PVD itself does share in the creative input. The business is described at the three levels of analysis in Figure 1.8.

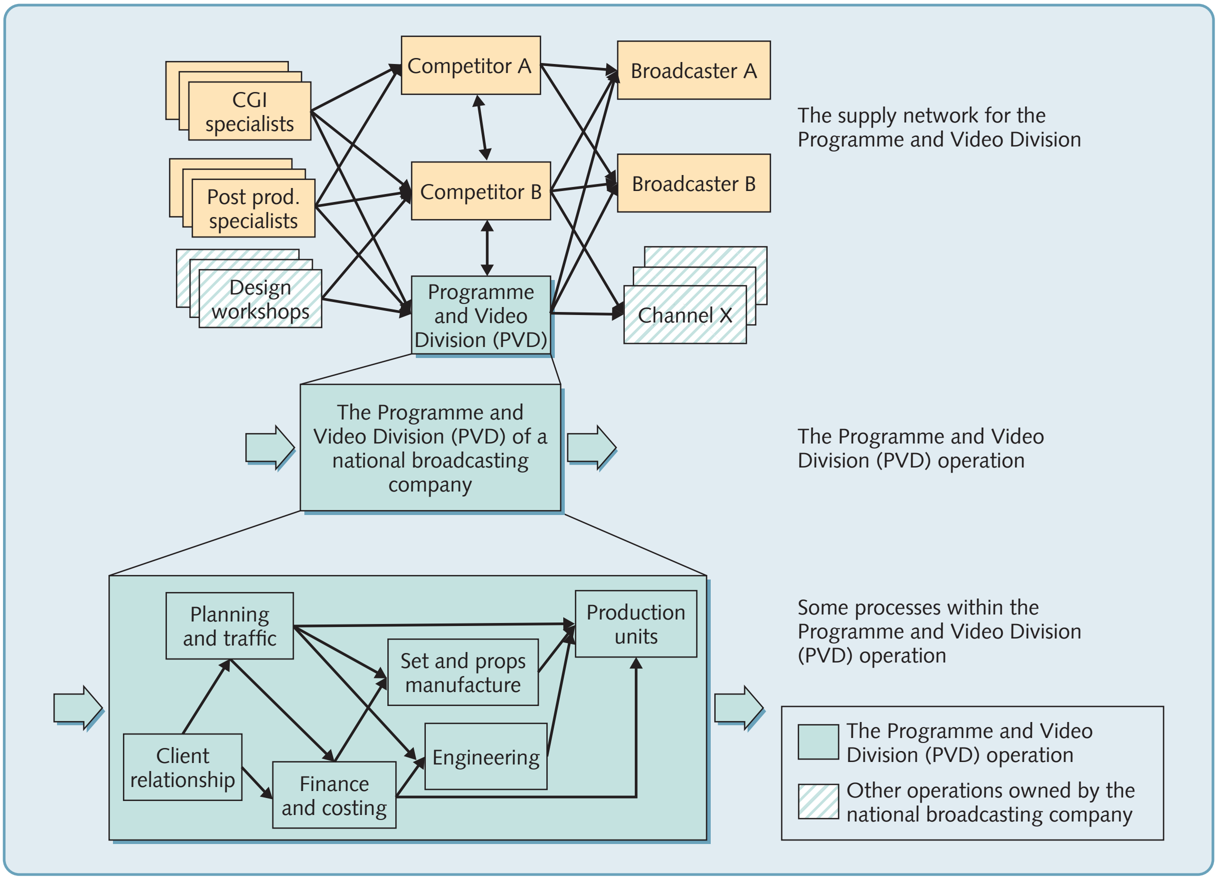

Figure 1.8 Operations and process management analysis for the programme and video division of a national broadcasting company (PVD) at three levels: the supply network, the operation and individual processes

At the level of the operation – the division produces products in the form of tapes, discs and media files, but its real 'product' is the creativity and 'artistry' that is captured in the programmes. *'We provide a service,'* said the division's boss, *'that interprets the client's needs (and sometimes their ideas), and transforms them into appealing and appropriate shows. We can do this because of the skills, experience and creativity of our staff, and our state-of-the-art technology.'*

At the level of the supply network – the division has positioned itself to specialise in certain types of product, including children's programmes, wildlife programmes and music videos. *'We did this so that we could develop a high level of expertise in a few relatively high margin areas. It also reduces our dependence on our own broadcasting channels. Having specialised in this way we are better positioned to partner and do work for other programme makers who are our competitors in some other markets. Specialisation has also allowed us to outsource some activities such as computer graphic imaging (CGI) and post-production that are no longer worth keeping in-house. However, our design workshop became so successful that they were "spun out" as a division in their own right and now work for other companies as well as ourselves.'*

At the level of individual processes – many smaller processes contribute directly or indirectly to the production of programmes and videos, including the following:

- The planning and traffic department who act as the operations managers for the whole operation; they draw up schedules, allocate resources and 'project manage' each job through to completion.

- Workshops that manufacture some of the sets, scenery and props for the productions.
- Client liaison staff who liaise with potential customers, test out programme ideas and give information and advice to programme makers.
- An engineering department that cares for, modifies and designs technical equipment.
- Production units that organise and shoot the programmes and videos.
- The finance and costing department that estimates the likely cost of future projects, controls operational budgets, pays bills and invoices customers.

Creating end-to-end processes – PVD produces several products and services that fulfil customer needs. Each of these, to different extents, involves several of the existing departments within the company. For example, preparing a 'pitch' (a sales presentation that includes estimates of the time and cost involved in potential projects) mainly needs the contributions of Client relations and the Finance and costing departments, but also needs smaller contributions from other departments. Figure 1.9 illustrates the contribution of each department to each product or service. (Figure 1.9 does not imply a particular sequence.) The contributions of each department may not all occur in the same order. Currently, all the division's processes are clustered into conventional departments defined by the type of activity they perform; engineering, client relationship, and so on. A radical redesign of the operation could involve regrouping activities and resources into five 'business' processes that fulfil each of the five defined customer

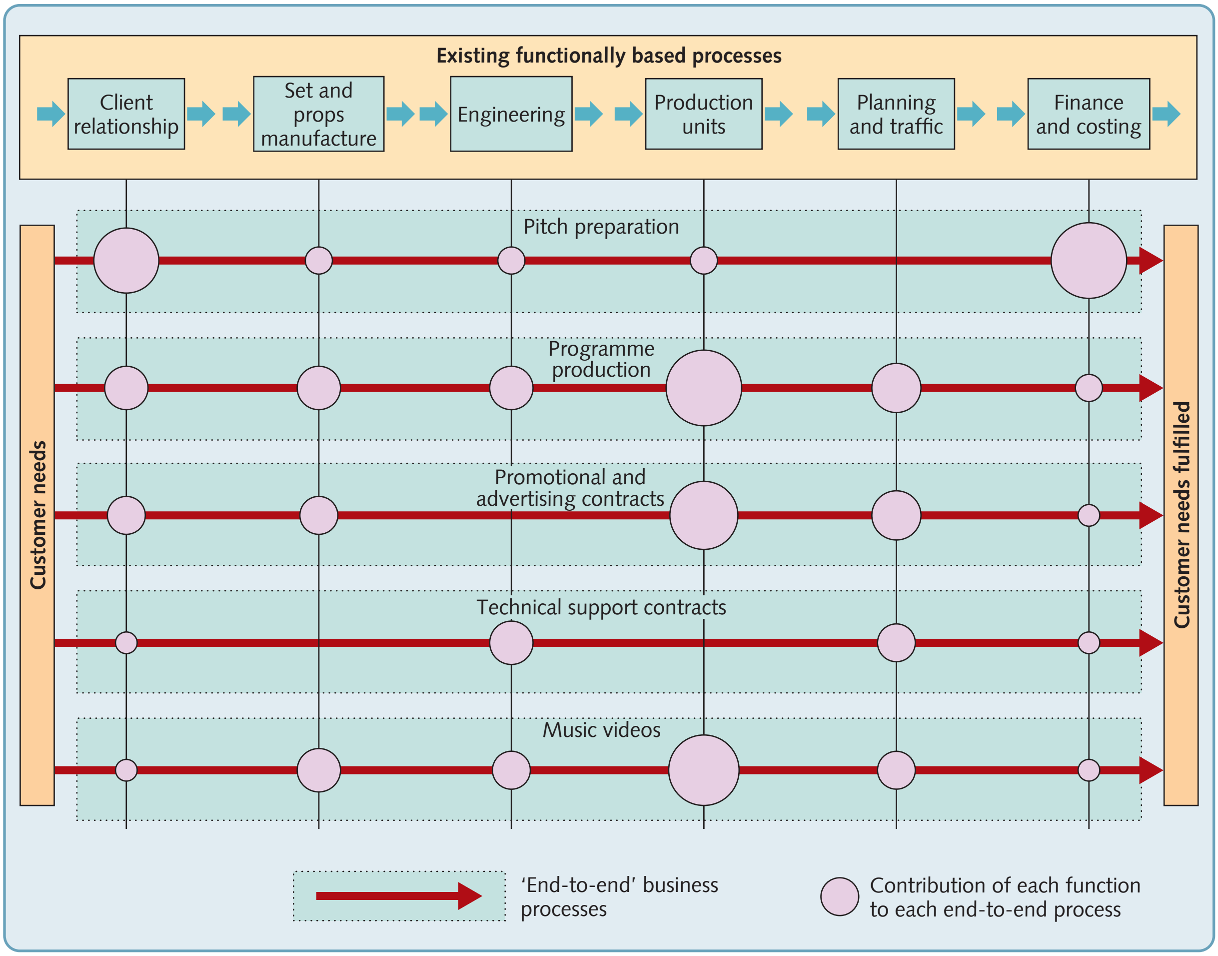

Figure 1.9 **An example of how processes in the Programme and Video Division (PVD) could be reorganised around end-to-end business processes that fulfil defined customer needs**

needs; this is shown diagrammatically by the dotted lines in Figure 1.9. It would involve the physical movement of resources (people and facilities) out of the current functional processes into the new end-to-end business processes. This is an example of how processes can be designed in ways that do not necessarily reflect conventional functional groupings.

DIAGNOSTIC QUESTION

Does operations and process management have a strategic impact?

One of the biggest mistakes a business can make is to confuse 'operations' with 'operation*al*'. Operational is the opposite of strategic; it means detailed, localised, short-term, day-to-day. 'Operations', on the other hand, is the set of resources that produce products and services. Operations can be treated both at an operational *and a strategic level.* We shall examine some views of operations strategy in the next chapter. For now, we treat a fundamental question for any operation – how does the way we manage operations and processes have a strategic impact?

The operations function of an organisation is clearly significant strategically, if only because it often represents the bulk of its assets and the majority of its people. Yet its true value is more than 'bulk'. It can 'make' the business in the sense that it gives the ability to compete through both the short-term ability to respond to customers and the long-term capabilities that will keep it ahead of its competitors. But if an operations function cannot produce its products and services effectively, it could 'break' the business by handicapping its performance no matter how it positions and sells itself in its markets. But in addition, all the processes within an organisation, irrespective of which function they support, can also help (or hinder) it in achieving its strategic objectives. It does this at three levels.

Three levels of operations performance

Operations performance can be assessed at different levels. At its broadest, operations can be judged on how it impacts on long-term societal issues. At the level of the individual firm or enterprise, it can be judged on how it supports (or not) their strategic aims. At its more operational level, the individual processes within an operation can be judged by how well they serve their (internal and external) customers and improve the efficiency with which they do it. In the rest of this section, we will look at each of these three levels.

Operations performance at a societal level

Operations decisions affect a whole variety of 'stakeholders' – the people and groups that have a legitimate interest in the operation's activities. These include the operation's employees, its customers, its shareholders, its immediate community and society in general. Although each of these groups will be interested in an operation's performance, they are likely to have very different views of which aspect of performance is important. Nevertheless, if one is to judge operations at a broad societal level, one must judge the impact it has on its stakeholders.

The triple bottom line

One idea that tries to capture the idea of a broader approach to assessing an organisation's performance is the 'triple bottom line'[4] (TBL, or 3BL), also known as 'people, planet and profit'. Essentially, it is a straightforward idea, simply that organisations should measure themselves not just on the traditional economic profit that they generate for their owners, but also on the impact their operations have on society (broadly, in the sense of communities, and individually, for example in terms of their employees) and the ecological impact on the environment. The influential initiative that has come out of this triple bottom line approach is that of 'sustainability'. A sustainable business is one that creates an acceptable profit for its owners, but minimises the damage to the environment and enhances the existence of the people with whom it has contact. In other words, it balances economic, environmental and societal interests. This gives the organisation its 'license to operate' in society. The assumption underlying the triple bottom line (which is not universally accepted, see the critical commentary at the end of the chapter) is that a sustainable business is more likely to remain successful in the long-term than one that focuses on economic goals alone.

OPERATIONS PRINCIPLE

Operations should judge themselves on the triple bottom line principle of, 'people, planet, profit'.

However, the idea of triple bottom line performance does have its critics. It is unreasonable, some say, for operations managers to have to cope with the conflicting pressures of maximising profitability, as well as managing in the interests of society in general with accountability and transparency. Even if a business wanted to reflect aspects of performance beyond its own immediate interests, how is it to do it? According to Michael Jensen of Harvard Business School, '*At the economy-wide or social level, the issue is this: If we could dictate the criterion or objective function to be maximized by firms (and thus the performance criterion by which corporate executives choose among alternative policy options), what would it be? Or, to put the issue even more simply: How do we want the firms in our economy to measure their own performance? How do we want them to determine what is better versus worse?*'[5] He also holds that using stakeholder perspectives gives undue weight to narrow special interests who want to use the organisation's resources for their own ends. The stakeholder perspective gives them a spurious legitimacy that '*undermines the foundations of value-seeking behaviour*'.

Operations performance at the level of the enterprise

The ability of operations and process management to impact the strategic success of any kind of enterprise is being increasingly recognised. When compared with only a few years ago, it attracts far more attention and, according to some reports, accounts for the largest share of all the money spent by businesses on consultancy advice. This may be partly because the area has been neglected in the past. But it also denotes an acceptance that it can have both short-term and long-term impact. This may be seen in the effect that operations and process management can have on the business's, cost, revenue, risk, investment, and capabilities.

- It can reduce the ***costs*** of producing products and services by being efficient. The more productive the operation is at transforming inputs into outputs, the lower will be the cost of producing a unit of output. Cost is never totally unimportant for any business, but generally the higher the cost of a product or service when compared to the price it commands in the market, the more important cost reduction will be as an operations objective. Even so, cost reduction is almost always treated as an important contribution that operations can make to the success of any business.
- It can increase ***revenue*** by increasing customer satisfaction through quality, service and innovation. Existing customers are more likely to be retained and new customers more likely to be attracted to products and services that are error-free and appropriately designed, if the operation is fast and responsive in meeting their needs and keeping its delivery promises, and if an operation can be flexible, both in customising its products and services and introducing new ones. It is operations that directly influence the quality, speed, dependability and flexibility of the business, all of which have a major impact on a company's ability to maximise its revenue.

- It can reduce the **risk** of operational failure, because well-designed and well-run operations should be less likely to fail. Eventually, all failures can be traced back to some kind of failure within a process. Furthermore, a well-designed process, if it does fail, should be able to recover faster and with less disruption (this is called *resilience*).
- It can ensure ***effective investment*** (*capital employed)* to produce its products and services. Eventually, all businesses in the commercial world are judged by the return that they produce for their investors. This is a function of profit (the difference between costs and revenues) and the amount of money invested in the business's operations resources. We have already established that effective and efficient operations can reduce costs and increase revenue. What is sometimes overlooked is operation's role in reducing the investment required per unit of output. It does this by increasing the effective capacity of the operation and by being innovative in how it uses its physical resources.
- It can ***build capabilities*** that will form the basis for *future* innovation by building a solid base of operations skills and knowledge within the business. Every time an operation produces a product or a service it has the opportunity to accumulate knowledge about how that product or service is best produced. If this accumulation of knowledge is used as a basis for learning and improvement, in the long term, capabilities can be built that will allow the operation to respond to future market challenges. Conversely, if an operations function is simply seen as the mechanical and routine fulfillment of customer requests, then it is difficult to build the knowledge base that will allow future innovation.

OPERATIONS PRINCIPLE

All operations should be expected to contribute to their enterprise by controlling costs, increasing revenue, reducing risks, making investment more effective and growing long-term capabilities.

Operations performance at the level of an operation's processes

Operations and process management can also be judged at the more operational level. This is how well the network of processes within an operation serves its internal, and eventually its external, customers. There are five aspects of operations and process performance, all of which to a greater or lesser extent will affect customer satisfaction and business competitiveness.

- **Quality** – doing things right, providing error-free goods and services that are 'fit for their purpose'.
- **Speed** – doing this fast, minimising the time between a customer asking for goods and services and the customer receiving them in full.
- **Dependability** – doing things on-time, keeping the delivery promises that have been made to customers.
- **Flexibility** – changing what you do or how you do it, the ability to vary or adapt the operation's activities to cope with unexpected circumstances or to give customers individual treatment, or to introduce new products or services.
- **Cost** – doing things cheaply, producing goods and services at a cost that enables them to be priced appropriately for the market, while still allowing a return to the organisation (or, in a not-for-profit organisation, that gives good value to the taxpayers or whoever is funding the operation).

OPERATIONS PRINCIPLE

Operations and process performance at an operational level can be grouped together as quality, speed, dependability, flexibility and cost.

However, do not think that these 'operational' aspects of performance have little strategic significance. On the contrary, they all contribute to the success of the organisation as a whole. We will look further at these aspects of performance in the next chapter that deals with operations strategy.

EXAMPLE

The programme and video division (PVD) continued

The PVD, described earlier, should be able to identify all four ways in which its operations and processes can have a strategic impact. It is expected the division will generate reasonable returns by controlling its costs and being able to command relatively high fees. *'Sure, we need to keep our costs down. We always review our budgets for bought-in materials and services. Just as important, we measure the efficiency of all our processes, and we expect annual improvements in process efficiency to compensate for any increases in input costs.* (Reducing costs). *Our services are in demand by customers because we are good to work with,'* said the division's Managing Director, *'We have the technical resources to do a really great job and we always give good service. Projects are completed on time and within budget. More importantly, our clients know that we can work with them to ensure a high level of programme creativity. That is why we can command reasonably high prices.'* (Increasing revenue). *'Also, we have a robust set of processes that minimize the chances of projects failing to achieve success.'* (Reducing risk). The division also has to justify its annual spend on equipment to its main board. *'We try and keep up to date with the new technology that can really make an impact on our programme making, but we always have to demonstrate how it will improve profitability.* (Effective investment.) *We also try to adapt new technology and integrate it into our creative processes in some way so that gives us some kind of advantage over our competitors.'* (Build capabilities.)

EXAMPLE

MSF operations provide medical aid to people in danger[6]

Do not assume that the operations and process perspective applies only to for-profit enterprises. Terms such as *competitive advantage, markets* and *business* used in this book are usually associated with companies in the for-profit sector. Yet, operations management is also relevant to organisations whose purpose is not primarily to earn profits. For example, Médecins Sans Frontières, MSF (Doctors Without Borders), the independent humanitarian organisation, uses its operations processes to provide medical aid where it is most needed in countries around the world, usually in crisis situations such as armed conflicts, epidemics, famines and natural disasters. Their teams deliver both medical aid and material aid quickly and efficiently. It is one of the most admired and effective relief organisations in the world. But no amount of fine intentions can translate into effective action without superior operations management. As MSF says, it must be able to react to any crisis with 'fast response, efficient logistics systems, and efficient project management'. Its response process has five phases: proposal, assessment, initiation, running the project and closing. The information that prompts a possible mission can come from governments, the international community, humanitarian organisations, or MSF teams already present in the region. Once the information has been checked and validated, MSF sends a team of medical and logistics experts to the crisis area to carry out a quick evaluation. The team assesses the situation, the number of people affected, and the current and future needs, and sends a proposal back to the MSF office. When the proposal is approved, MSF staff start the process of selecting personnel, organising materials and resources and securing project funds. Initiating a project involves sending technical equipment and resources to the area. In large crises, planes fly in all the necessary materials so that the work can begin immediately. Thanks to their pre-planned processes, specialised kits and the emergency stores, MSF can distribute material and equipment within 48 hours, ready for the response team to start work as soon as they arrive. Most MSF projects generally run for somewhere between 18 months and three and a half years. Whether an emergency response or a long-term healthcare project, the closing process is roughly similar. Once the critical medical needs have been met (which could be after weeks, months or years depending on the situation), MSF begins to close the project with a gradual withdrawal of staff and equipment. At this stage, the project closes or is passed onto an appropriate organisation. MSF will also close a project if risks in the area become too great to ensure staff safety.

DIAGNOSTIC QUESTION

Are processes managed to reflect their operating circumstances?

All processes differ in some way, so, to some extent, all processes will need to be managed differently. Some of the differences between processes are 'technical' in the sense that different products and services require different skills and technologies to produce them. However, processes also differ in terms of the nature of demand for their products or services. Four characteristics of demand, in particular, have a significant effect on how processes need to be managed:

- The volume of the products and services produced.
- The variety of the different products and services produced.
- The variation in the demand for products and services.
- The degree of visibility that customers have of the production of products and services.

> OPERATIONS PRINCIPLE
> *The way in which processes need to be managed is influenced by volume, variety, variation and visibility.*

Volume – Processes with a high volume of output will have a high degree of repeatability, and because tasks are repeated frequently it often makes sense for staff to specialise in the tasks they perform. This allows the systemisation of activities, where standard procedures may be codified and set down in a manual with instructions on how each part of the job should be performed. Because tasks are systemised and repeated, it is often worthwhile developing specialised technology that gives higher processing efficiencies. By contrast, low-volume processes with less repetition cannot specialise to the same degree. Staff are likely to perform a wide range of tasks, and while this may be more rewarding, it is less open to systemisation. Nor is it likely that efficient, high-throughput technology could be used. The implications of this are that high-volume processes have more opportunities to produce products or services at low unit cost. For example, the volume and standardisation of large fast-food restaurant chains such as McDonalds or KFC, enables them to produce with greater efficiency than a small, local cafeteria or diner.

Variety – Processes that produce a high variety of products and services must engage in a wide range of different activities, changing relatively frequently between each activity. It must also contain a wide range of skills and technology, sufficiently 'general purpose' to cope with the range of activities and sufficiently flexible to change between them. A high level of variety may also imply a relatively wide range of inputs to the process and the additional complexity of matching customer requirements to appropriate products or services. High-variety processes are invariably more complex and costly than low-variety ones. For example, a taxi company is usually prepared to pick up and drive customers almost anywhere (at a price); they may even take you by the route of your choice. There are an infinite number of potential routes (products) that it offers. But, its cost per kilometre travelled will be higher than a less customised form of transport such as a bus service.

Variation – Processes are generally easier to manage when they only have to cope with predictably constant demand. Resources can be geared to a level that is just capable of meeting demand. All activities can be planned in advance. By contrast, when demand is variable and/or unpredictable, resources will have to be adjusted over time. Worse still, when demand is unpredictable, extra resources will have to be designed into the process to provide a 'capacity cushion' that can absorb unexpected demand. For example, processes that manufacture high-fashion garments will have to cope with the general seasonality of the garment market, together with the uncertainty of whether particular styles may or may not prove popular. Operations that make conventional business suits are likely to have less fluctuation in demand over time, and be less prone to unexpected fluctuations. Because processes with lower variation do not need any extra safety capacity and can be planned in advance, they will generally have lower costs than those with higher variation.

Visibility – Process visibility is a slightly more difficult concept to envisage. It indicates how much of the processes are 'experienced' directly by customers, or how much the process is 'exposed' to its customers. Generally, processes that act directly on customers (such as retail processes or health care process) will have more of their activities visible to their customers than those that act on materials and information. However, even material- and information-transforming processes may provide a degree of visibility to the customers. For example, parcel distribution operations provide internet-based 'track and trace' facilities to enable their customers to have visibility of where their packages are at any time. Low-visibility processes, if they communicate with their customers at all, do so using less immediate channels such as the telephone or the internet. Much of the process can be more 'factory-like'. The time lag between customer request and response could be measured in days rather than the near-immediate response expected from high-visibility processes. This lag allows the activities in a low-visibility process to be performed when it is convenient to the operation, so achieving high utilisation. Also, because the customer interface needs managing, staff in high-visibility processes need customer contact skills that shape the customer's perception of process performance. For all these reasons, high-visibility processes tend to have higher costs than low-visibility processes.

Many operations have both high- and low-visibility processes. This serves to emphasise the difference that the degree of visibility makes. For example, in an airport, some of its processes are relatively visible to its customers (check in desks, the information desks, restaurants, passport control and security staff, etc.). These staff operate in a high-visibility 'front-office' environment. Other processes in the airport have relatively little, if any, customer visibility (baggage handling processes, overnight freight operations, loading meals on to the aircraft, cleaning, etc). We rarely see these processes; they perform the vital but low-visibility tasks, in the 'back-office' part of the operation.

The implications of the four Vs of processes

All four dimensions have implications for processing costs. Put simply, high volume, low variety, low variation and low visibility all help to keep processing costs down. Conversely, low volume, high variety, high variation and high customer contact generally carry some kind of cost penalty for the process. This is why the volume dimension is drawn with its 'low' end at the left, unlike the other dimensions, to keep all the 'low-cost' implications on the right. Figure 1.10 summarises the implications of such positioning. The four dimensions also have implications for the type of activities that operations managers will have to focus on. As we will see in later chapters, volume and variety are particularly important in how processes are resourced and designed (Chapters 5 and 6), variation in demand calls for careful consideration of how capacity is to be managed (Chapter 8), while visibility has a significant influence over how the quality of products and services are judged (Chapter 13).

Charting processes using the 4 Vs

In almost any operation, processes can be identified that have different positions on the four dimensions, and which therefore have different objectives and will need managing in different ways. To a large extent the position of a process on the four dimensions is determined by the demand of the market it is serving. However, most processes have some discretion

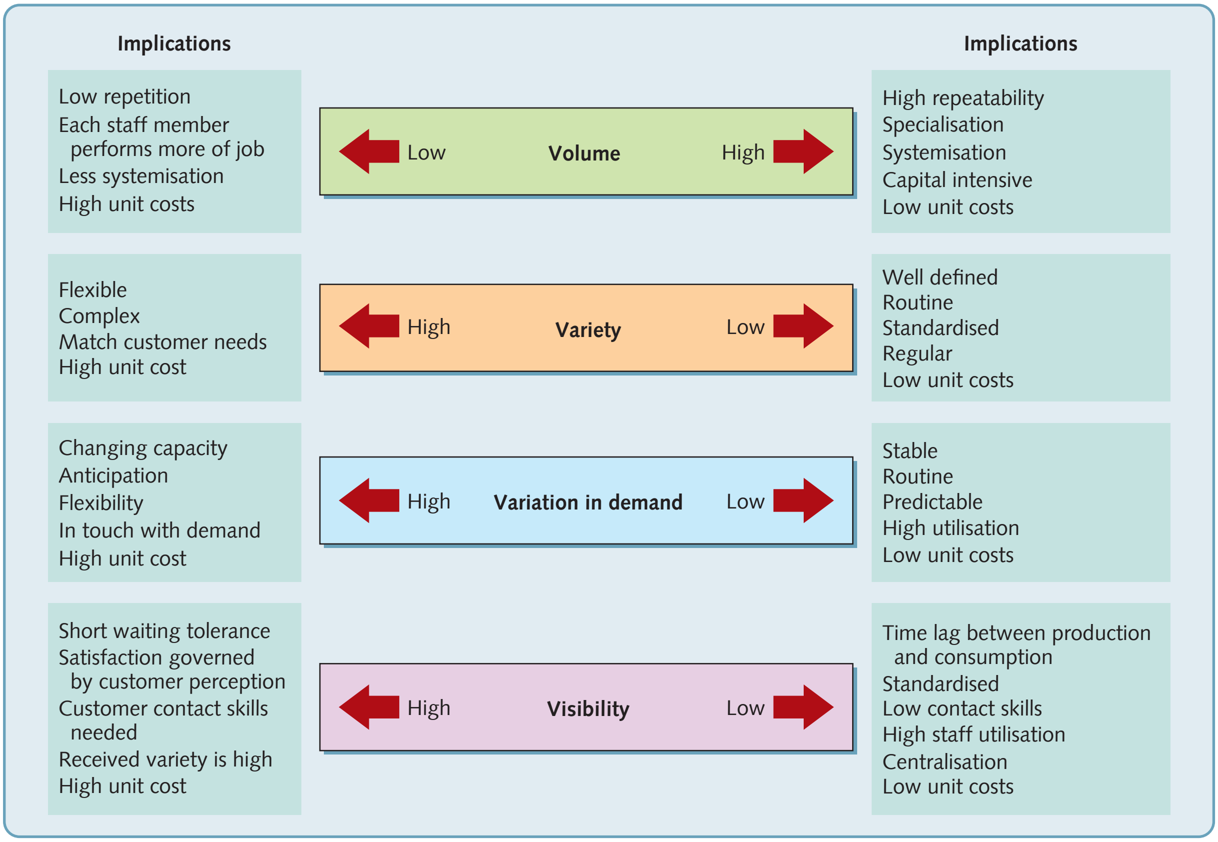

Figure 1.10 A typology of operations

in moving themselves on the dimensions. Look at the different positions on the visibility dimension that retail banks have adopted. At one time, using branch tellers was the only way customers could contact a bank. Now access to the bank's services could be through (in decreasing order of visibility) a personal banker, who visits your home or office, a conversation with a branch manager, the teller at the window, telephone contact through a call centre, internet banking services, or an ATM cash machine. These other processes offer services that have been developed by banks to serve different market needs.

> **OPERATIONS PRINCIPLE**
> *Operations and processes can (other things being equal) reduce their costs by increasing volume, reducing variety, reducing variation and reducing visibility.*

Figure 1.11 illustrates the different positions on the four Vs for some retail banking processes. Note that the personal banking/advice service is positioned at the high-cost end of the four Vs. For this reason, such services are often only offered to relatively wealthy customers that represent high-profit opportunities for the bank. Note also that the more recent developments in retail banking such as call centres, internet banking and ATMs, all represent a shift towards the low-cost end of the four Vs. New processes that exploit new technologies can often have a profound impact on the implications of each dimension. For example, internet banking, when compared with an ATM cash machine, offers a far higher variety of options for customers but because the process is automated through its information technology, the cost of offering this variety is less than at a conventional branch or even a call centre.

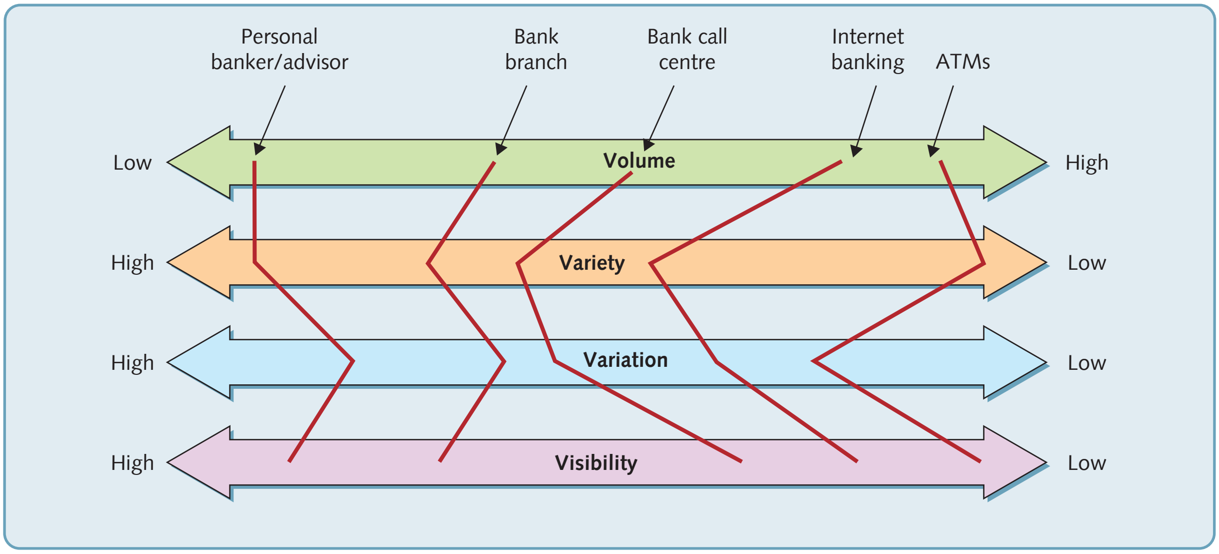

Figure 1.11 Four Vs analysis for some retail banking processes

DIAGNOSTIC QUESTION

Are operations and process decision-making appropriate?

> **OPERATIONS PRINCIPLE**
>
> *Operations management activities can be grouped into four broad categories: directing the overall strategy of the operation; designing the operation's products, services and processes, planning and controlling delivery; and developing process performance.*

Managing operations and processes involves a whole range of separate decisions that will determine how well they achieve their overall purpose and contribute to the organisation as a whole. The way operations managers approach decision-making is therefore of considerable importance in determining their effectiveness. But different operations decisions serve different purposes. They can be grouped together in various ways. Look at other books on operations management and you will find many different ways of structuring operations decisions and therefore the subject as a whole. Here we have chosen to classify activities into four broad groups, relating to four broad activities that, more or less, follow a sequence that corresponds to the life cycle of operations and processes:

1. ***Directing*** the overall strategy of the operation. A general understanding of operations and processes and their strategic purpose, together with an appreciation of how strategic purpose is translated into reality through how innovation is incorporated into products and services and how much of the total value-adding process should be kept in-house and how much outsourced.
2. ***Designing*** the operation's processes. Design is the activity of determining the physical form, shape and composition of operations and processes, together with the type of resources they contain.
3. Planning and control process ***delivery.*** After being designed, the delivery of products and services from suppliers and through the total operation to customers must be planned and controlled.
4. ***Developing*** process performance. Increasingly, it is recognised that operations, or any process managers cannot simply routinely deliver products and services in the same way

that they always have done. They have a responsibility to develop the capabilities of their processes to improve process performance.

A model of operations and process management

We can now combine two ideas to develop the model of operations and process management that will be used throughout this book. The first is the idea that *operations* and the *processes* that make up both the operations and other business functions are transformation systems that take in inputs and use process resources to transform them into outputs. The second idea is that the resources, both in an organisation's operations as a whole and in its individual processes, need to be managed in terms of how they are *directed,* how they are *designed,* how *delivery* is planned and controlled and how they are *developed* and improved. Figure 1.12 shows how these two ideas go together. This book will use this classification to examine the more important decisions that should be of interest to all managers of operations and processes.

However, although it is useful to categorise operations decisions in this way, the boundaries between these four categories are not 'clean'. There are always overlaps and interrelationships between the categories. The decisions that we classify as 'directing' will, of course, impact on all other operations decisions – that is the definition of 'directing'. And how processes are designed will limit how delivery can be organised and how easy it is to develop their capabilities. 'Delivery' decisions, such as whether to adopt lean practices, very much impact on how improvement (development) happens, and so on. All operations decisions are interrelated. Our categorisation of decisions simply indicates their main purpose.

Quantitative and qualitative decision-making in operations

In operations and process management (and throughout this book) we frequently use 'models', to better understand a decision. By a 'model' we mean an explicit statement of our image of reality. It represents those aspects of a decision with which we are concerned. It structures and

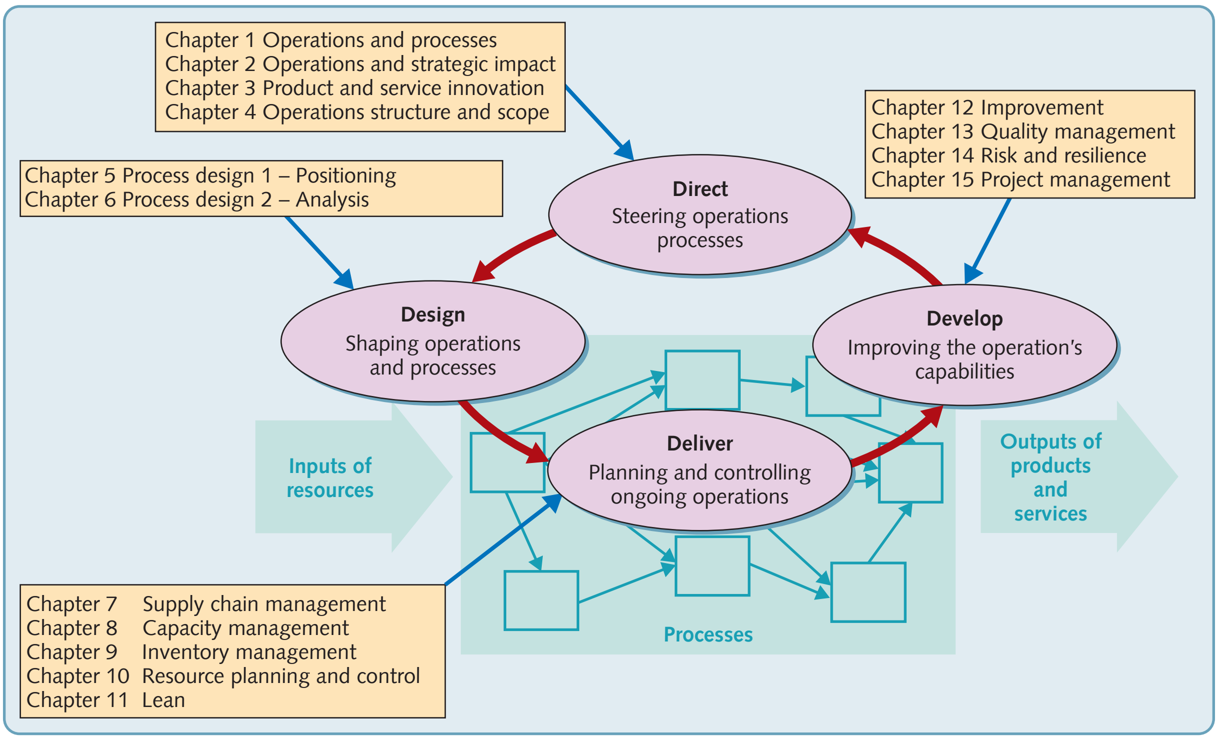

Figure 1.12 **Operations and processes management: a general model**

formalises the information we possess about the decision and, in doing so, presents reality in a simplified and organised form. A model, therefore provides an abstraction of a more complex reality. They can be partial in that they exclude some factors, and can aggregate or compress several factors into one, but models are at the core of operations decision-making. Some of the models used in operations are qualitative. They categorise or describe the relationships between aspects of decisions, but do not necessarily ascribe precise associations between variables. For example, Figure 1.8, which describes the relationships between processes and operations for a programme and video division of a national broadcasting company, is essentially a qualitative model. It does not provide an 'answer' for any decision, but it does enhance understanding and stimulate discussion around other possible ways to organise the operations.

Quantitative models are also important in operations and process management, but present different challenges. Quantitative models try to represent the underlying behaviours involved in a decision by using mathematical and/or statistical descriptions of relationships. They allocate numerical values to variables to produce a mathematic representation of reality. For example, the economic order quantity (EOQ) model that we explain in Chapter 9 is a good illustration of a quantitative model. It gives a precise relationship between the costs involved in making one particular inventory decision and therefore can be used to make the decision of how much stock to order. Well, at least it is supposed to. In fact, this model illustrated one of the problems with using a quantitative approach in operations and process management. In order to model the decision mathematically, reality has to be simplified to an extent that may severely limit its usefulness. Not that this is a condemnation of quantitative modelling. Practical operations management depends on the quantification of decision-making where possible. But for most operations decisions, some combination of quantitative and qualitative modelling is required.

'Behavioural' operations

Academics who write about, research, or teach, operations management are sometimes accused of ignoring the 'practical reality' of how operations management decisions are made. It is claimed that their models, frameworks and guidance do not reflect how people really behave when making operations management decisions in practice. This has led to the development of a (relatively) new branch of operations management. It is called 'behavioural operations management' (BOM), or simply 'behavioural operations', and explores the 'interaction of human behaviours and operational systems and processes'[7]. More specifically, it challenges the idea that managers are rational when making decisions that impact operations performance. One research team summarises what they see as common behavioural assumptions to operations models[8]:

- People are not a major factor in operations decisions.
- People are deterministic and predictable.
- People make decisions independently of each other.
- People do not learn from experience.
- People are not part of the product or service.
- People are emotionless.
- Work is perfectly observable and can be understood.

Clearly, these assumptions are extremely unrealistic, and while no experienced operations manager would ever subscribe to them, they do act as a warning as to how the models, frameworks and techniques in this book should *not* be interpreted. Very few of the models that we use are rigidly prescriptive. Generally, they do not attempt to dictate a single 'optimum' solution to any problem. Rather they try to structure and clarify operations decisions, the better to understand, debate and, hopefully make, better decisions.

Critical commentary

All chapters will contain a short critical commentary on the main ideas covered in the chapter. Its purpose is not to undermine the issues discussed in the chapter, but to emphasise that, although we present a relatively orthodox view of operation, there are other perspectives.

- The central idea in this introductory chapter is that all organisations have operations (and other functions) that have processes that produce products and services, and that all these processes are essentially similar. However, some believe that by even trying to characterise organisations in this way (perhaps even by calling them 'processes') one loses or distorts their nature and depersonalises or takes the 'humanity' out of the way in which we think of the organisation. This point is often raised in not-for-profit organisations, especially by 'professional' staff. For example, the head of one European 'Medical Association' (a Doctors' Trade Union) criticised hospital authorities for expecting a 'sausage factory service based on productivity targets'. No matter how similar they appear on paper, it is argued, a hospital can never be viewed in the same way as a factory. Even in commercial businesses, professionals, such as creative staff, often express discomfort at their expertise being described as a 'process'.

- At least some of this discomfort at the idea that the process perspective is demeaning in some way is the (false) assumption that 'people' and 'process' are different things. 'Which is more important to your business', some critics say, 'people or processes?' This implies that people are somehow not part of processes. A process is 'an arrangement of resources and activities that transform inputs into outputs', where one of the most important resources is usually the people that staff it, and the activities are carried out (largely, if not exclusively) by people. People and processes are not different things. People are integral to processes and processes are what people do.

- However, to some extent the criticisms of taking an exclusively process (or mechanistic) perspective are valid. How we describe organisations does say much about our underlying assumptions of what an 'organisation' is and how it is supposed to work. Notwithstanding the point we made earlier about how a purely process view can misleadingly imply that organisations are neat and controllable with unambiguous boundaries and lines of accountability, a process perspective can risk depicting the messy reality of organisations in a naïve manner. Yet, in our view it is a risk well worth taking.

SUMMARY CHECKLIST

Each chapter contains a summary checklist in the form of questions that can be usefully applied to any type of operations and reflect the major diagnostic questions used within the chapter.

- ☐ Is the operations function of the business clearly defined?
- ☐ Do operations managers realise that they are operations managers even if they are called by some other title?
- ☐ Do the non-operations functions within the business realise that they manage processes?
- ☐ Does everyone understand the inputs, activities and outputs of the processes of which they are part?
- ☐ Is the balance between products and services produced by the operations function well understood?
- ☐ Are future changes that may occur in the balance between products and services produced by the operation understood?
- ☐ Does the operation assess its impact on the environment and society as well as its financial performance?
- ☐ What contribution is operations making towards reducing the cost of products and services?
- ☐ What contribution is operations making towards increasing the revenue from products and services?
- ☐ What contribution is operations making towards reducing the risks of failure and increasing the effectiveness of recovery?
- ☐ What contribution is operations making towards better use of capital employed?
- ☐ How is operations developing the capability for future innovation?
- ☐ Does the operation understand its position in the overall supply network?
- ☐ Does the operation contribute to the overall supply network?
- ☐ Are the individual processes that comprise the operations function defined and understood?
- ☐ Are individual processes aware of the internal customer and supplier concept?
- ☐ Do they use the internal customer and supplier concept to increase their contribution to the business as a whole?
- ☐ Do they use the ideas and principles of operations management to improve the performance of their processes?
- ☐ Has the concept of end-to-end business processes been examined and considered?
- ☐ Are the differences (in terms of volume, variety, variation, and visibility) between processes understood?
- ☐ Are the volume, variety, variation and visibility characteristics of processes reflected in the way they are managed?
- ☐ Do all operations managers understand the full range of decisions that they should be involved in?
- ☐ Are all operations managers aware of the various qualitative and quantitative models that can help them in their decision-making?

CASE STUDY

Design house partnerships at concept design services

'I can't believe how much we have changed in a relatively short time. From being an inward-looking manufacturer, we became a customer-focused "design and make" operation. Now we are an "integrated service provider". Most of our new business comes from the partnerships we have formed with design houses. In effect, we design products jointly with specialist design houses that have a well-known brand, and offer them a complete service of manufacturing and distribution. In many ways we are now a "business-to-business" company rather than a "business-to-consumer" company.' (Jim Thompson, CEO, Concept Design Services (CDS))

Concept Design Services (CDS) had become one of Europe's most profitable homeware businesses. It had moved in two stages from making precision plastic components, mainly in the Aerospace sector, together with some cheap 'homeware' items such as buckets and dustpans, sold under the 'Focus' brand name, to making very high quality (expensive) stylish homewares with a high 'design value' for well- known brands.

The first stage – from 'Focus' to 'Concept'

The initial move into higher margin homeware had been masterminded by Linda Fleet, CDS's Marketing Director. 'My previous experience in the decorative products industry had taught me the importance of fashion and product development, even in mundane products such as paint. Premium-priced colours and new textures would catch the popular imagination and need supporting by promotion and editorial features in lifestyle magazines. The players who embraced this fashion element of the market were dramatically more profitable than those who simply provided standard ranges. Instinctively, I felt that this must also apply to homeware. We decided to develop a whole coordinated range of such items, and to open up a new distribution network for them to serve the more exclusive stores, kitchen equipment and specialty retailers. Within a year of launching our first new range of kitchen homeware under the 'Concept' brand name, we had over 3,000 retail outlets across Northern Europe with full point-of-sale display facilities and supported by press coverage and product placement on TV 'life style' programmes. Within two years 'concept' products were providing over 75 per cent of our revenue and 90 per cent of our profits'. (The margin on Concept products is many times higher than for the Focus range.) During this period the Focus (basic) range continued to be produced, but as a drastically reduced range.

The second stage – from 'Concept' to 'Design House partnerships'

Linda was also the driving force behind the move to design house partnerships. *'It started as a simple design collaboration between our design team and an Italian "design house".* (Design houses are creative product designers who may, or may not own a brand of their own, but rarely manufacture or distribute their products, relying on outsourcing to subcontractors.) *It seemed a natural progression to them asking us to first manufacture and then distribute this and other of their designs. Over the next five years, we built up this business, so now we design (often jointly with the design house), manufacture and distribute products for several of the more prestigious European design houses. We think this sort of business is likely to grow. The design houses appreciate our ability to offer a full service. We can design products in conjunction with their own design staff and offer them a level of manufacturing expertise they can't get elsewhere. More significantly, we can offer a distribution service which is tailored to their needs. From the customer's point of view the distribution arrangements appear to belong to the design house itself. In fact they are based exclusively on our own call centre, warehouse and distribution resources.'*

The most successful collaboration was with Villessi, the Italian designers. Generally, it was CDS's design expertise that was attractive to 'design house' partners. Not only did CDS employ professionally respected designers, they had also acquired a reputation for being able to translate difficult technical designs into manufacturable and saleable products. Design house partnerships usually involved relatively long lead times but produced unique products with very high margins, nearly always carrying the design house's brand.

Manufacturing operations

All manufacturing was carried out in a facility located 20 km from Head Office. Its moulding area housed large injection-moulding machines; most with robotic material handling capabilities. Products and components passed to the packing hall, where they were assembled and inspected. The newer, more complex products often had to move from moulding to assembly and then back again for further moulding. All products followed the same broad process route but, with more products needing several progressive moulding and assembly stages, there was an increase in 'process flow recycling', which was adding complexity. One idea was to devote a separate cell to the newer and more complex products until they had 'bedded in'. This cell could also be used for testing new moulds. However, it would need investment in extra capacity that would not always be fully utilised. After manufacture, products were packed and stored in the adjacent distribution centre.

'When we moved into making the higher margin "concept" products, we disposed of most of our older, small injection-moulding machines. Having all larger machines allowed us to use large multi-cavity moulds. This increased productivity by allowing us to produce several products, or components, each machine cycle. It also allowed us to use high-quality and complex moulds, which, although cumbersome and more difficult to changeover, were very efficient and gave a very high-quality product. For example, with the same labour we could make 3 items per minute on the old machines, and 18 items per minute on the modern ones using multi moulds. That's a 600 per cent increase in productivity. We also achieved high dimensional accuracy, excellent surface finish, and extreme consistency of colour. We could do this because of our expertise derived from years of making aerospace products. By standardising on single large machines, any mould could fit any machine. This was an ideal situation from a planning perspective, as we were often asked to make small runs of Concept products at short notice.' (Grant Williams, CDS Operations Manager)

Increasing volume and a desire to reduce cost had resulted in CDS subcontracting much (but not all) of its Focus products to other (usually smaller) moulding companies. *'We would never do it with any complex or Design House partner products, but it should allow us to reduce the cost of making basic products, while releasing capacity for higher margin ones. However, there have been quite a few "teething problems". Coordinating the production schedules is currently a problem, as is agreeing quality standards. To some extent, it's our own fault. We didn't realise that subcontracting was a skill in its own right. And although we have got over some of the problems, we still do not have a satisfactory relationship with all of our subcontractors.'* (Grant Williams, CDS Operations Manager)

Planning and distribution services

The distribution services department of the company was regarded as being at the heart of the company's customer service drive. Its purpose was to integrate the efforts of design, manufacturing and sales by planning the flow of products from production, through the distribution centre to the customer. Sandra White, the Planning Manager, reported to Linda Fleet and was responsible for the scheduling of all manufacturing and distribution, and for maintaining inventory levels for all the warehoused items. '*We try to stick to a preferred production sequence for each machine and mould, so as to minimise set-up times, by starting on a light colour and progressing through a sequence to the darkest. We can change colours in 15 minutes, but because our moulds are large and technically complex, mould changes can take up to three hours. Good scheduling is important to maintain high plant utilisation. With a higher variety of complex products, batch sizes have reduced and it has brought down average utilisation. Often we can't stick to schedules. Short-term changes are inevitable in a fashion market. Certainly better forecasts would help. . . but even our own promotions are sometimes organised at such short notice that we often get caught with stockouts. New products, in particular, are difficult to forecast, especially when they are "fashion" items and/or seasonal. Also, I have to schedule production time for new product mould trials; we normally allow 24 hours for the testing of each new mould received, and this has to be done on production machines. Even if we have urgent orders, the needs of the designers always have priority.'* (Sandra White)

Customer orders for Concept and Design House partnership products were taken by the company's sales call centre, located next to the warehouse. The individual orders would then be dispatched using the company's own fleet of medium and small distribution vehicles for UK orders, but using carriers for the Continental European market. A standard delivery timetable was used and an 'express delivery' service was offered to those customers prepared to pay a small delivery premium. However, a recent study had shown that almost 40 per cent of express deliveries were initiated by the company, rather than customers. Typically, this would be to fulfil deliveries of orders containing products that were out of stock at the time of ordering. The express delivery service was not required for *Focus* products because almost all deliveries were made to five large customers. The size of each order was usually very large, with deliveries to customers' own distribution depots. However, although the organisation of Focus delivery was relatively straightforward, the consequences of failure were large. Missing a delivery meant upsetting a large customer.

Challenges for CDS

Although the company was financially successful and very well regarded in the homeware industry, there were a number of issues and challenges that it knew it would have to address. The first was the role of the design department and its influence over new product development. New product development had become particularly important to CDS, especially since they had formed alliances with design houses. This had led to substantial growth in both the size and the influence of the design department, which reported to Linda Fleet. *'Building up and retaining design expertise will be the key to our future. Most of our growth is going to come from the business which will be brought in through the creativity and flair of our designers. Those who can combine creativity with an understanding of our partners' business and design needs can now bring in substantial contracts. The existing business is important of course, but growth will come directly from these people's capabilities.'* (Linda Fleet)

But not everyone was so sanguine about the rise of the Design Department. *'It is undeniable that relationships between the designers and other parts of the company have been under strain recently. I suppose it is, to some extent, inevitable. After all, they really do need the freedom to design as they wish. I can understand it when they get frustrated at some of the constraints that we have to work under in the manufacturing or distribution parts of the business. They also should be able to expect a professional level of service from us. Yet, the truth is that they make most of the problems themselves. They sometimes don't seem to understand the consequences or implications of their design decisions or the promises they make to the design houses. More seriously, they don't really understand that we could actually help them do their job better if they cooperated a bit more. In fact, I now see some of our Design House partners' designers more than I do our own designers. The Villessi designers are always in my factory and we have developed some really good relationships.' (Grant Williams)*

The second major issue concerned sales forecasting, and again there were two different views. Grant Williams was convinced that forecasts should be improved. *'Every Friday morning we devise a schedule of production and distribution for the following week. Yet, usually before Tuesday morning, it has had to be significantly changed because of unexpected orders coming in from our customers' weekend sales. This causes tremendous disruption to both manufacturing and distribution operations. If sales could be forecast more accurately, we would achieve far high utilization, better customer service, and I believe, significant cost savings.'*

However, Linda Fleet saw things differently. *'Look, I do understand Grant's frustration, but after all, this is a fashion business. By definition, it is impossible to forecast accurately. In terms of month-by-month sales volumes, we are in fact pretty accurate, but trying to make a forecast for every week end every product is almost impossible to do accurately. Sorry, that's just the nature of the business we're in. In fact, although Grant complains about our lack of forecast accuracy, he always does a great job in responding to unexpected customer demand.'*

Jim Thompson, the Managing Director, summed up his view of the current situation. *'Particularly significant has been our alliances with the Italian and German design houses. In effect, we are positioning ourselves as a complete service partner to the designers. We have a world-class design capability, together with manufacturing, order processing, order-taking and distribution services. These abilities allow us to develop genuinely equal partnerships which integrate us into the whole industry's activities'*.

Linda Fleet also saw an increasing role for collaborative arrangements. *'It may be that we are seeing a fundamental change in how we do business within our industry. We have always seen ourselves as primarily a company that satisfies consumer desires through the medium of providing good service to retailers. The new partnership arrangements put us more into the "business- to-business" sector. I don't have any problem with this in principle, but I'm a little anxious as to how much it gets us into areas of business beyond our core expertise.'*

The final issue which was being debated within the company was longer term, and particularly important. *'The two big changes we have made in this company have both happened because we exploited a strength we already had within the company. Moving into Concept products was only possible because we brought our high-tech precision expertise that we had developed in the aerospace sector into the homeware sector, where none of our new competitors could match our manufacturing excellence. Then, when we moved into Design House partnerships, we did so because we had a set of designers who could command respect from the world class design houses with whom we formed partnerships. So what is the next move for us? Do we expand globally? We are strong in Europe but nowhere else in the world. Do we extend our design scope into other markets, such as furniture? If so, that would take us into areas where we have no manufacturing expertise. We are great at plastic injection moulding, but if we tried any other manufacturing processes, we would be no better than, and probably worse than, other firms with more experience. So what's the future for us?'* (Jim Thompson, CEO CDS).

QUESTIONS

1. Why is operations management important in CDS?
2. Draw a 4Vs profile for the company's products/services
3. What would you recommend to the company if they asked you to advise them in improving their operations?

APPLYING THE PRINCIPLES

Some of these exercises can be answered by reading the chapter. Others will require some general knowledge of business activity and some might require an element of investigation. **Model answers for the first two questions are available on the student companion website.** For model answers to the other questions in this section, please ask your tutor/lecturer.

1 Quentin Cakes make about 20,000 cakes per year in two sizes, both based on the same recipe. Sales peak at Christmas time when demand is about 50 per cent higher than in the more quiet summer period. Their customers (the stores who stock their products) order their cakes in advance through a simple internet-based ordering system. Knowing that they have some surplus capacity, one of their customers has approached them with two potential new orders.

The *Custom Cake* Option – this would involve making cakes in different sizes where consumers could specify a message or greeting to be 'iced' on top of the cake. The consumer would give the inscription to the store who would e-mail it through to the factory. The customer thought that demand would be around 1,000 cakes per year, mostly at celebration times such as Valentine's Day and Christmas.

The *Individual Cake* Option – this option involves Quentin Cakes introducing a new line of about 10-15 types of very small cakes intended for individual consumption. Demand for this individual-sized cake was forecast to be around 4,000 per year, with demand likely to be more evenly distributed throughout the year than their existing products.

The total revenue from both options is likely to be roughly the same and the company has only capacity to adopt one of the ideas. Which one should it be?

2 Described as having '*revolutionised the concept of sandwich making and eating*', Prêt A Manger opened their first shop in the mid-1980s, in London. Now they have over 130 shops in the UK, New York, Hong Kong and Tokyo. They say that their secret is to focus continually on quality, in all its activities. '*Many food retailers focus on extending the shelf life of their food, but that's of no interest to us. We maintain our edge by selling food that simply can't be beaten for freshness. At the end of the day, we give whatever we haven't sold to charity to help feed those who would otherwise go hungry.* The first Prêt A Manger shop had its own kitchen, where fresh ingredients were delivered first thing every morning, and food was prepared throughout the day. Every Prêt shop since has followed this model. The team members serving on the tills at lunchtime will have been making sandwiches in the kitchen that morning. They rejected the idea of a huge centralised sandwich factory, even though it could significantly reduce costs. Prêt also own and manage all their shops directly so that they can ensure consistently high standards. '*We are determined never to forget that our hardworking people make all the difference. They are our heart and soul. When they care, our business is sound. If they cease to care, our business goes down the drain. We work hard at building great teams. We take our reward schemes and career opportunities very seriously. We don't work nights (generally), we wear jeans, we party!*'

(a) Do you think Prêt A Manger fully understand the importance of their operations management?

(b) What evidence is there for this?

(c) What kind of operations management activities at Prêt A Manger might come under the four headings of direct, design, deliver and develop?

3 Visit an IKEA superstore and a smaller furniture store. Observe how the shop operates, for example, where customers go, how staff interact with them, how big it is, how the shop has chosen to use its space, what variety of products it offers, and so on. Talk with the staff and managers if you can. Think about how the two shops differ from each other. Then consider the question, '*What implications do the differences between IKEA and the smaller shop have for their operations management?*'

4 Write down five services that you have 'consumed' in the last week. Try to make these as varied as possible. Examples could include public transport, a bank, any shop or supermarket, attendance at an education course, a cinema, a restaurant, and so on.
For each of these services, ask yourself the following questions.

- Did the service meet your expectations? Is so what did the management of the service have to do well in order to satisfy your expectations? If not, where did they fail? Why might they have failed?
- If you were in charge of managing the delivery of these services, what would you do to improve the service?
- If they wanted to, how could the service be delivered at a lower cost so that the service could reduce its prices?
- How do you think that the service copes when something goes wrong (such as a piece of technology breaking down)?
- Which other organisations might supply the service with products and services? (In other words, they are your 'supplier', but who are *their* suppliers)?
- How do you think the service copes with fluctuation of demand over the day, week, month or year?

These questions are just some of the issues that the operations managers in these services have to deal with. Think about the other issues they will have to manage, in order to deliver the service effectively.

5 Find a copy of a financial newspaper (*Financial Times, Wall Street Journal, Economist,* etc.) and identify one company that is described in the paper that day. What do you think would be the main operations issues for that company?

Notes on chapter

1 Sources include: Trangbæk, R. R. (2015) 'Lego group among five most reputable companies in the world', LEGO Group Press office, 21 April; Diaz, J. (2014) 'Exclusive look inside the Lego Factory', Gizmodo .com, http://lego.gizmodo.com/exclusive-look-inside-the-lego-factory-5022769; Terdiman, D. (2011) 'Watching Lego make its world-famous bricks', CNET, http://www.cnet.com/uk/news/watching-lego-make-its-world-famous-bricks/, 15 June; *The Economist* (2014) 'Unpacking Lego - How the Danish firm became the world's hottest toy company', 8 March.

2 Sources include discussions with Tom Dyson and Olly Willans of Torchbox, to whom we are grateful for their advice and assistance.

3 Vargo, S.L. and Lusch, R.F. (2008), 'Service-dominant logic: Continuing the evolution', *Journal of the Academy of Marketing Science,* 36(Spring), 1–10.

4 The phrase 'the triple bottom line' was first used in 1994 by John Elkington, the founder of a British consultancy called SustainAbility. Read Elkington, J. (1997) *Cannibals with Forks: the Triple Bottom Line of 21st Century Business,* Capstone. Also good is, Savitz, A.W. and Weber, K. (2006) *The Triple Bottom Line: How Today's Best-Run Companies Are Achieving Economic, Social and Environmental Success—and How You Can Too,* Jossey-Bass.

5 Jensen, M. C. (2001) 'Value maximization, stakeholder theory, and the corporate objective function', *Journal of Applied Corporate Finance,* vol 14, no 3, 7–21.

6 Sources include: MSF website (2014) http://www.msf.org.uk; Peter Beaumont (2011) 'Médecins sans Frontières book reveals aid agencies' ugly compromises', The *Guardian,* 20 November.

7 Eckerd, S. and Elliot Bendoly (2010) 'Behavioral operations', *Scholarpedia,* 5(3):10422.

8 Boudreau, J., Hopp, W., McClain, J.O., and Thomas, L.J. (2003) 'On the interface between operations and human resources management', *Manufacturing & Service Operations Management,* Vol. 26, no. 2, 179–202.

TAKING IT FURTHER

Anupindi, R. and Chopra, S. (2013) ***Managing Business Process Flows (3rd edn), Pearson.*** *Takes a 'process' view of operations, it's mathematical but rewarding.*

Barnes, D. (2007) ***Operations Management: An international perspective, Cengage Learning.*** *A text that is similar in outlook to this one, but with more of a (useful) international perspective.*

Hall, J.M. and Johnson, M.E. (2009) 'When Should a Process Be Art, Not Science?', ***Harvard Business Review, March.*** *One of the few articles that looks at the boundaries of conventional process theory.*

Johnston, R., Clark, E. and Shulver, M. (2012) ***Service Operations Management*** (4th edn), ***Pearson.*** *What can we say! A great treatment of service operations from the same stable as this textbook.*

Lewis, M.A. and Brown, A.D. (2011) 'How different is professional service operations management?' ***Journal of Operations Management, vol. 30, no. 1.*** *It's an academic article, but it does say something interesting about professional service operations.*

Schmenner, R. W., Van Wassenhove, L., Ketokivi, M., Heyl, J. and Lusch, R.F. (2009) 'Too much theory, not enough understanding', ***Journal of Operations Management, vol. 27, 339–343.*** *An academic treatment of one aspect of the operations management literature.*

Slack, N. (2017) ***The Operations Advantage, Kogan Page.*** *Similar coverage to this book (and just as brilliant), but aimed very much at a practitioner audience.*

Slack, N. and Lewis, M. (eds) (2005) ***The Blackwell Encyclopedic Dictionary of Operations Management, (2nd edn) Blackwell Business, Oxford.*** *For those who like technical descriptions and definitions.*

2

Operations and strategic impact

Introduction

In the long-term, the major (and some would say, only) objective for operations and processes is to provide a business with some form of strategic advantage. That is the reason why the management of a business's processes and operations and its intended overall strategy must be logically connected. Yet for many in business, the very idea of an 'operations strategy' is a contradiction in terms. After all, to be involved in the strategy process is the complete opposite of those day-to-day tasks and activities associated with being an operations manager. Nevertheless, it is also clear that operations can have a real strategic impact. Many *enduringly* remarkable enterprises, from Apple to Zara use their operations resources to gain long-term strategic success. Such firms have found that it is the way they manage their operations that sets them apart from, and above, their competitors. But if operations strategy can be the making of business success, a failure to have one, or a failure to adapt it to circumstances, can break a business. Without a strong link with overall strategy, operations and processes will be without a coherent direction. And without direction they may finish up making internal decisions that either do not reflect strategy, or that conflict with each other, or both. So, although operations and process management is largely 'operational', it also has a strategic dimension that is vital if operations is to fulfill its potential to contribute to competitiveness. Figure 2.1 shows the position of the ideas described in this chapter in the general model of operations management.

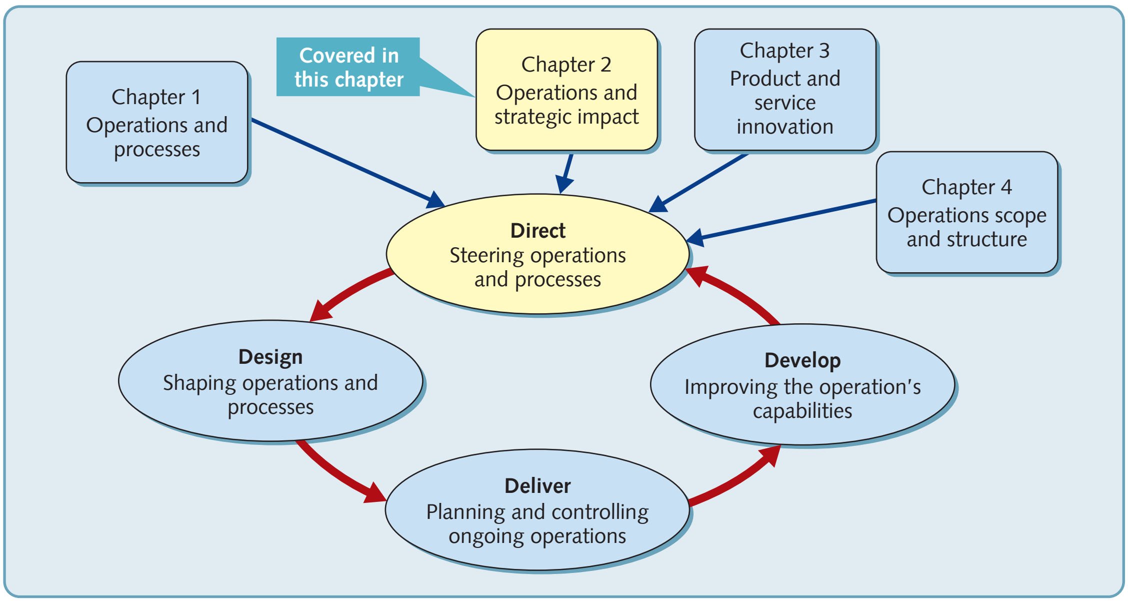

Figure 2.1 Operations strategy is the pattern of decisions and actions that shape the long-term vision, objectives and capabilities of the operation and its contribution to overall strategy

EXECUTIVE SUMMARY

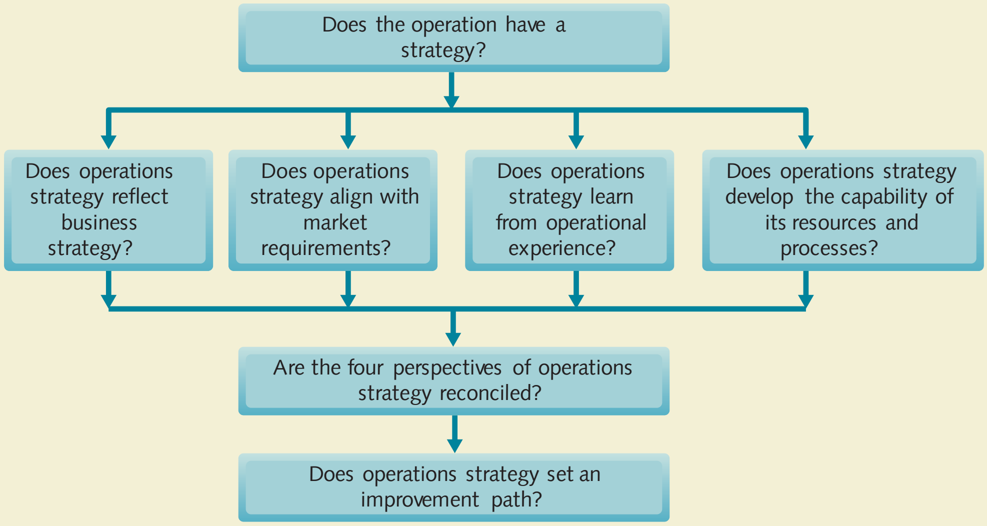

Does operations have a strategy?

Operations strategy is the pattern of decisions and actions that shape the long-term vision, objectives and capabilities of the operation and its contribution to overall strategy. It is the way in which operations resources are developed over the long term to create sustainable competitive advantage for the business. Increasingly, many businesses are seeing their operations strategy as one of the best ways to differentiate themselves from competitors. Even in those companies that are marketing-led (such as fast-moving consumer goods), an effective operations strategy can add value by allowing the exploitation of market positioning. Strategies are always difficult to identify because they have no presence in themselves, but are identified by the pattern of decisions that they generate. Nevertheless, one can identify what an operations strategy should do. First, it should take significant stakeholders into account. They are the people and groups who have a legitimate interest in the operation's strategy. They include employees, customers, society or community groups, shareholders, suppliers, and industry regulators. Second, it should articulate a vision for the operations contribution. This is a clear statement of how operations intend to contribute value for the business. A common approach to summarising operations contribution is the Hayes and Wheelwright Four-Stage Model.

Does operations strategy reflect business strategy (top-down)?

A top-down perspective often identifies three related levels of strategy: corporate strategy that should position the corporation in its global, economic, political and social environment; business strategy that sets out each business's individual mission and objectives; and functional strategy that considers what part each function should play in contributing to the strategic objectives of the business. Any top-down perspective on operations strategy should attempt to achieve both correspondence (it is consistent with business strategy) and coherence (its various elements all 'pull' in the same direction). One way of achieving this is to devise a 'business model' that integrates with an 'operating model'.

Does operations strategy align with market requirements (outside-in)?

Operations strategy should reflect the market position of the business. And because companies compete in different ways, the operations function should respond by providing the ability to perform in a manner that is appropriate for the intended market position. This is called a market (or outside-in) perspective on operations strategy. It involves translating market requirements into 'performance objectives' (quality, speed, dependability, flexibility and cost), the exact definition of which will be different for different operations. These may be refined by distinguishing between those performance objectives that are 'order-winners' (those things that directly and significantly contribute to winning business), and those that are 'qualifiers' (the aspects of competitiveness where the operation's performance has to be above a particular level just to be considered by the customer).

Does operations strategy learn from operational experience (bottom-up)?

Operations may adopt a particular strategic direction not because of any formal high-level decision-making, but because the ongoing experience of providing products and services to customers at an operational level convinces them that it is the right thing to do. This is the 'bottom-up' perspective of operations strategy. It stresses the concept of 'emergent strategies' where strategy-making can be relatively unstructured and fragmented. The important attributes for shaping strategy from the bottom up are an ability to capture the learning that should come from routine, operations activities and being able to transform that learning into strategically valuable knowledge.

Does operations strategy develop the capability of its resources and processes (inside-out)?

But this is not the only objective. In the longer term, operations strategy must build the capabilities within its resources and processes that will allow the business to provide something to the market that its competitors find difficult to imitate or match. The idea of the basis of long-term competitive capabilities deriving from the operation's resources and processes is called the resource, or inside-out, perspective on operations strategy. It is very much influenced by the resource-based view (RBV) of the firm. 'Difficult-to-imitate' resources can be classified as 'strategic' if they are scarce, or imperfectly mobile, or imperfectly imitable, or imperfectly substitutable.

Are the four perspectives of operations strategy reconciled?

None of the four perspectives individually can give a full picture of any operations strategy. But together they do provide a good idea of how its operations are contributing strategically. Yet, the four perspectives must be reconciled. One way to do this is to construct an operations strategy matrix. This is a simple model that describes operations strategy as the intersection of a company's performance objectives and the strategic decisions that it makes. Alternatively, the 'line of fit' model can be used to map the balance between requirements of the market and the capabilities of the operation over time. Ideally, there should be a reasonable degree of alignment, or 'fit' between the requirements of the market and the capabilities of the operation.

Does operations strategy set an improvement path?

The purpose of operations strategy is to improve the business's performance relative to its competitors' in the long-term. It therefore must provide an indication of how this improvement is to take place. This is best addressed by considering the trade-offs between performance objectives in terms of the 'efficient frontier' model. This describes operations strategy as a combination of repositioning performance along an existing efficient frontier, and increasing overall operations effectiveness by overcoming trade-offs to expand the efficient frontier.

DIAGNOSTIC QUESTION

Does the operation have a strategy?

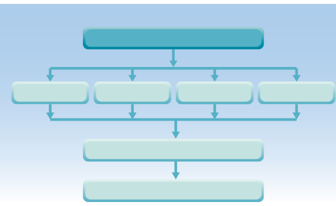

Let us start with the meaning of 'strategy'. Surprisingly, is not particularly easy to define. Linguistically the word derives from the Greek word '*strategos*' meaning 'leading an army'. And although there is no direct historical link between Greek military practice and modern ideas of strategy, the military metaphor is powerful. Both military and business strategy can be described in similar ways, and include some of the following:

- Setting broad objectives that direct an enterprise towards its overall goal.
- Planning the path (in general rather than specific terms) that will achieve these goals.
- Stressing long-term rather than short-term objectives.
- Dealing with the total picture rather than stressing individual activities.
- Being detached from, and above, the confusion and distractions of day-to-day activities.

All these points apply to operations strategy. It is the 'pattern of decisions and actions that shape the long-term vision, objectives and capabilities of the operation and its contribution to the overall strategy of the business'.[1] Yet, the term 'operations strategy' sounds at first like a contradiction. How can 'operations', a subject that is generally concerned with the day-to-day creation and delivery of goods and services, be strategic? 'Strategy' is usually regarded as the opposite of those day-to-day routine activities. But, as we indicated previously, '*operations*' is not the same as '*operational'*. 'Operations' are the resources and processes that create products and services. 'Operational' is the opposite of strategic. It means day-to-day, detailed and often localised.

Perhaps more significantly, many of the businesses that seem to be especially successful, and that appear to be sustaining their success into the longer term, have a clear and often inventive operations strategy. Just look at some of the high-profile companies quoted in this book, or that feature in the business press. It is not just that their operations strategy provides these companies with adequate support; it is their operations strategy that is the pivotal reason for their competitive superiority. Just as revealing, when companies stumble, it is often because they have either taken their eye off the operations ball, or failed to appreciate its importance in the first place. Operations strategy can make or break a business. More generally, all enterprises, *and all parts of the enterprise,* need to prevent strategic decisions being frustrated by poor operational implementation. And this idea leads us to the second purpose of this chapter (and indeed the book as a whole): to show how, by using the principles of operations strategy, *all* parts of any business, and *all* functions of a business can contribute effectively to the overall success of the business. So the idea of 'operations strategy' has two different but related meanings. The first is concerned with the operations function itself, and how it can contribute to strategic success. The second is concerned with how *any* function can develop its resources and processes to establish its strategic role.

EXAMPLE

Amazon

As a publicly stated ambitious target it takes some beating. 'Amazon.com strives to be', it says, 'Earth's most customer-centric company'. Founded by Jeff Bezos in 1995, the Amazon.com website started as a place to buy books, giving its customers what at the time was a unique customer experience. Bezos believed that only the internet could offer customers the convenience of browsing a selection of millions of book titles in a single sitting. During its first 30 days of business, Amazon.com fulfilled orders for customers in 45 countries – all shipped from Bezos's Seattle-area garage. And that initial success has been followed by continued growth that is based on a clear strategy of technological innovation. Among its many technological innovations for customers, Amazon.com offers a personalised shopping experience for each customer, book discovery through 'Search Inside the Book', convenient checkout using '1-Click® Shopping' and community features like Listmania and Wish Lists that help customers discover new products and make informed buying decisions. In addition, Amazon operates retail websites and offers programs that enable other retailers and individual sellers to sell products on their websites. It may not be glamorous, but Amazon has focused on what have been called 'the dull-but-difficult tasks' such as tracking products, managing suppliers, storing inventory and delivering boxes. Fulfilment By Amazon allows other companies to use Amazon's logistics capability including the handling of returned items, and access to Amazon's 'back-end' technology.

Amazon Web Services' cloud computing business, provides the computing power for small and larger high-profile customers such as Spotify's digital music service, and Netflix's video streaming service. But, why should any business want to allow Amazon to have such control over its activities? Mainly because it allows entrepreneurs to create start-ups and established companies to expand their activities without the huge investment they would need to build appropriate infrastructure themselves. Amazon's large and efficient operations are also better value than smaller companies could achieve. Now, many prominent retailers work with Amazon Services to power their e-commerce offerings from end-to-end, including technology services, merchandising, customer service and order fulfilment. Offering business-to-business services is also good for Amazon. The problem with on-line retailing, said Bezos, is its seasonality. At peak times, such as Christmas, Amazon has far more computing capacity than it needs for the rest of the year. At low points it may be using as little as 10 per cent of its total capacity. Hiring out that spare capacity is an obvious way to bring in extra revenue. Its EC2 (Elastic Compute Cloud) service provides resizable computing capacity 'in the cloud'. It is designed, said Amazon, to make web-scale computing easier for developers . . . *Amazon EC2's simple web service interface allows you to obtain and configure capacity with minimal friction. It provides you with complete control of your computing resources and lets you run on Amazon's proven computing environment. Amazon EC2 reduces the time required to obtain and boot new server instances to minutes, allowing you to quickly scale capacity, both up and down, as your computing requirements change. Amazon EC2 changes the economics of computing by allowing you to pay only for capacity that you actually use. Amazon EC2 provides developers with the tools to build failure-resilient applications and isolate themselves from common failure scenarios.* Don't worry if you can't follow the technicalities of Amazon's statement; it is aimed at IT professionals. The important point is that it is a business-to-business service based on the company's core competence of leveraging its processes and technology that can make retail operations ultra-efficient. However, some observers immediately criticised Amazon's apparent redefinition of its strategy. 'Why not,' they said, 'stick to what you know, focus on your core competence of internet retailing?' Bezos's response was clear. 'We *are* sticking to our core competence. The only thing that's changed is that we are exposing it for (the benefit of) others.' At least for Jeff Bezos, Amazon is not so much an internet retailer as a provider of internet-based technology and logistics services.

EXAMPLE

SSTL

Sometimes a firm's operations strategy can change an industry. Look at space satellites. They can be very expensive. And in the early days of space missions, only superpowers could afford to develop and launch them. The conventional wisdom was that space was such a hostile environment that satellites would have to be constructed using only expensive specially developed components. Yet, in the late 1970s this assumption was challenged by Sir Martin Sweeting, who then was studying for his PhD at the University of Surrey in the UK. The aerospace research team in the Electrical Engineering Department at the University of Surrey had built its first satellite using commercial off-the-shelf components. It was about as big as two microwave ovens, weighing in at 72 kg (by contrast, some of the satellites being launched by government space agencies were as large as a London double-decker bus). It was launched in 1981 with the help of NASA. The team followed this up in 1984 with a second satellite that they built in six months. A year later Surrey Satellite Technology Limited (SSTL) was formed. The firm's vision was to open up the market for space exploration by pioneering the use of small and relatively cheap, but reliable, satellites built from readily available off-the-shelf components. It was a revolutionary idea. Now SSLT is the world's leading small satellite company, and has delivered space missions for a whole range of applications including Earth observation, science, communications and in-orbit technology demonstration. The company is at the forefront of space innovation, exploiting advances in technologies and challenging conventions to bring affordable space exploration to international customers. The company, which employs more than 500 staff, has launched over 40 satellites and is based across four sites in South East England. Since 2014 SSTL has been an independent company within the Airbus Defence and space group.

How has SSTL achieved this success from such small beginnings? Partly because it was an early player in the market, having the vision to see that there would be a market for small satellites that could serve the ambitions of smaller countries, companies, research groups and even schools. But in addition, it has always been innovative in finding ways of keeping the cost of building the satellites down. SSTL pioneered the low-cost, low-risk approach to delivering operational satellite missions within short development timescales and with the capability that potential customers wanted. Particularly important was the company's use of commercial off-the-shelf technology. In effect, using industry-standard parts meant exploiting the, often enormous, investments by consumer-electronics companies, auto part manufacturers, and others who had developed complex components for their products. Even if this sometimes limited what a satellite could do, it provided the scale economies that would be impossible if they were designing and making customised components from scratch. 'We were being parasitic, if you like,' admits Sir Martin.

However, not all commercially available components made for terrestrial use are up to coping with conditions in space, which is a hugely important issue. Reliability is essential in a satellite. (It's difficult to repair them once in space.) And even though off-the-shelf components and systems have become increasing reliable, they must be rigorously tested to make sure that they are up to the severe conditions found in space. Knowing which bits can be used and which cannot is an important piece of knowledge. Yet, although individual components and systems are often bought off-the-shelf, the company does most of its operations activities itself. This allows SSTL to provide a complete in-house design, manufacture, launch and operation service as well as a range of advice, analysis and consultancy services. 'What distinguishes us is our vertically integrated capability, from design and research to manufacturing and operations,' said Sir Martin. 'We don't have to rely on suppliers, although of course we buy in components when that is advantageous.' And innovation? It's still as important as it was at the company's start.

What do these two examples have in common?

Both examples exhibited several important characteristics that one finds in many successful companies. First, both were innovative in how they have exploited the potential of emerging technologies. Amazon used its operations strategy to take advantage of the potential released by technological change in how we could order goods, how they could be delivered, how data from transactions could be exploited and how the resources that were developed to achieve this could be further exploited (cloud computing services). SSTL realised that, with clever design, relatively inexpensive components could achieve a large part of what far more costly components could. Second, both linked the development of their operations to a well-defined idea of what their customers wanted. All aspects of their operations strategy had clear customer benefits. Third, both companies have learned from their experience. Given that it processes millions of transactions every week, Amazon also takes the opportunity to learn from its everyday operations (which is why it is able to predict what you might like to buy). While a key part of SSTL's core knowledge is which components can stand up to conditions in space. Finally, and most significantly, both companies actually did have an operations strategy; both realised the importance of strategically directing their operations resources. Neither company ignored the strategic potential of their operations resources and processes.

How do you judge an operations strategy?

There are some problems in asking this apparently simple question. In most operations management decisions you can see what you are dealing with. You can touch inventory, talk to people, programme machines, and so on. But strategy is different. You cannot see a strategy, feel it, or touch it. Also, whereas the effects of most operations management decisions become evident relatively fast, it may be years before an operations strategy decision can be judged to be a success or not. Moreover, any 'strategy' is always more than a single decision. Operations strategy will be revealed in the total *pattern* of decisions that a business takes in developing its operations in the long term. Nevertheless, the question is an obvious starting point and one that must be addressed by all operations.

An operations strategy should take significant stakeholders into account

OPERATIONS PRINCIPLE

Operations strategy should take significant stakeholders into account.

All operations have stakeholders. They are the people and groups who have a legitimate interest in the operation's strategy. Some are internal (employees); others are external (customers, society or community groups, and a company's shareholders). External stakeholders may have a direct commercial relationship with the organisation (suppliers and customers); others may not (industry regulators). In not-for-profit operations, these stakeholder groups can overlap. So, voluntary workers in a charity may be employees, shareholders and customers all at once. However, in any kind of organisation, it is a responsibility of the operations function to understand the (often conflicting) objectives of its stakeholders and set its objectives accordingly. Yet, although all stakeholder groups, to different extents, will be interested in operations performance, they are likely to have very different views on which aspect of performance is important. Table 2.1 identifies typical stakeholder requirements. But stakeholder relationships are not just one way. It is also useful to consider what an individual organisation or business wants of the stakeholder groups themselves. Some of these requirements are illustrated in Table 2.1. The dilemma with using this wide range of stakeholders to judge performance is that organisations, particularly commercial companies, have to cope with the conflicting pressures of maximising profitability on the one hand, with the expectation that they will manage in the interests of (all or part of) society in general with accountability and transparency. Even if a business wanted to reflect aspects of performance beyond its own immediate interests, how is it to do it?

Table 2.1 Typical stakeholders' performance objectives

Stakeholder	*What stakeholders want from the operation*	*What the operation wants from stakeholders*
Shareholders	Return on investment Stability of earnings Liquidity of investment	Investment capital Long-term commitment
Directors/top management	Low/acceptable operating costs Secure revenue Well-targeted investment Low risk of failure Future innovation	Coherent, consistent, clear and achievable strategies Appropriate investment
Staff	Fair wages Good working conditions Safe work environment Personal and career development	Attendance Diligence/best efforts Honesty Engagement
Staff representative bodies (e.g. Trade Unions)	Conformance with national agreements Consultation	Understanding Fairness Assistance in problem solving
Suppliers (of materials, services, equipment, etc.)	Early notice of requirements Long-term orders Fair price On-time payment	Integrity of delivery, quality and volume Innovation Responsiveness Progressive price reductions
Regulators (e.g. Financial regulators)	Conformance to regulations Feedback on effectiveness of regulations	Consistency of regulation Consistency of application of regulations Responsiveness to industry concerns
Government (Local, National, Regional)	Conformance to legal requirements Contribution to (local/national/regional) economy	Low/simple taxation Representation of local concerns Appropriate infrastructure
Lobby groups (e.g. Environmental Lobby Groups)	Alignment of the organisation's activities with whatever the group are promoting	No unfair targeting Practical help in achieving aims (if the organisation wants to achieve them)
Society	Minimise negative effects from the operation (noise, traffic, etc.) and maximise positive effects (jobs, local sponsorship, etc.	Support for organisation's plans

Corporate social responsibility (CSR)

Strongly related to the stakeholder perspective of operations performance is that of corporate social responsibility (generally known as CSR). A direct link with the stakeholder concept is to be found in the definition used by Marks and Spencer, the UK-based retailer. '*Corporate Social Responsibility . . . is listening and responding to the needs of a company's stakeholders. This includes the requirements of sustainable development. We believe that building good relationships with employees, suppliers and wider society is the best guarantee of long-term success. This is the backbone of our approach to CSR.*' The issue of how broader social performance objectives can be included in operations management's activities is of increasing importance, both from an ethical and a commercial point of view. However, converting the CSR concept into operational reality presents considerable difficulty, although several attempts have been made. For example, the 'triple bottom line' approach that we described in the previous chapter is one of the best-known attempts to integrate economic, environmental and social impacts.

An operations strategy should articulate a vision for the operations contribution

> **OPERATIONS PRINCIPLE**
> ***Operations strategy should articulate a 'vision' for the operations function contribution to overall strategy.***

The 'vision' for an operation is a clear statement of how operations intend to contribute value for the business. It is not a statement of what the operation wants to *achieve* (those are its objectives), but rather an idea of what it must *become* and what contribution it should make. A common approach to summarising operations contribution is the Hayes and Wheelwright Four-Stage Model.[2] The model traces the progression of the operations function from what is the largely negative role of 'Stage 1' operations to it becoming the central element of competitive strategy in excellent 'Stage 4' operations. Figure 2.2 illustrates the four steps involved in moving from Stage 1 to Stage 4

Stage 1: Internal neutrality

This is the very poorest level of contribution by the operations function. The other functions regard it as holding them back from competing effectively. The operations function is inward-looking and at best reactive, with very little positive to contribute towards competitive success. Its goal is to be ignored. At least then it isn't holding the company back in any way. Certainly, the rest of the organisation would not look to operations as the source of any originality, flair or competitive drive. Its vision is to be 'internally neutral', a position it attempts to achieve not by anything positive but by avoiding the bigger mistakes.

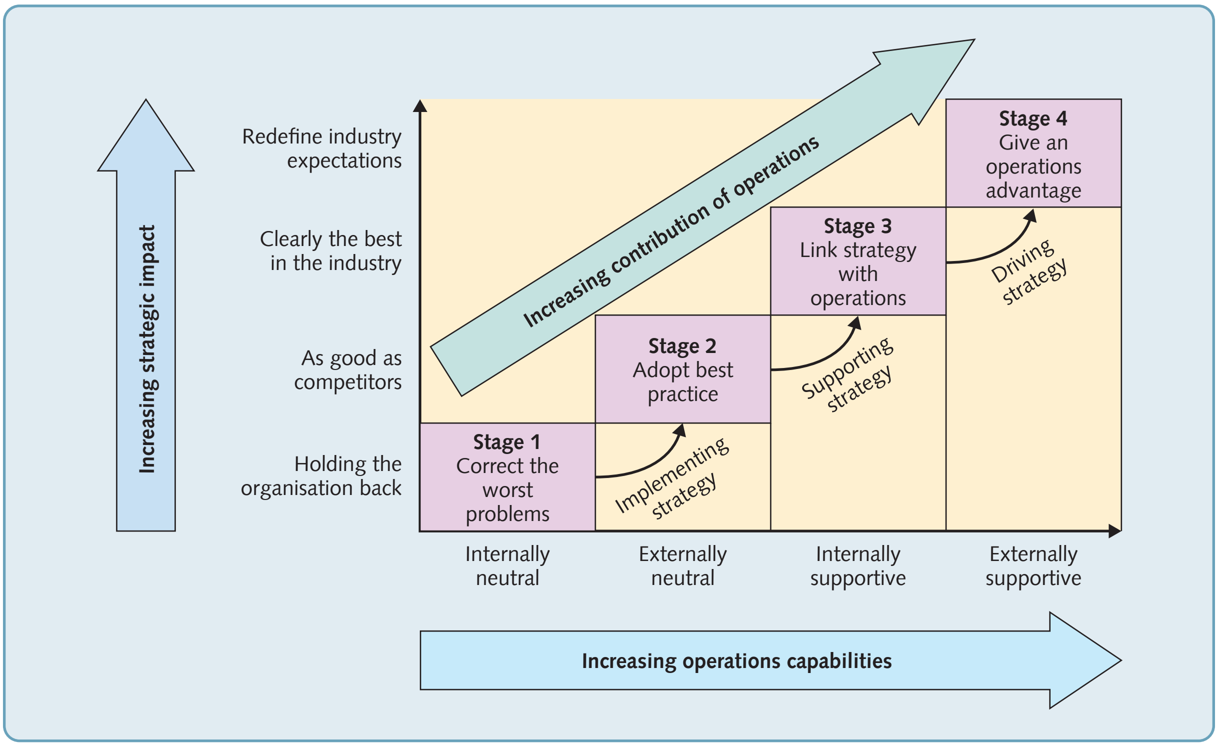

Figure 2.2 **Hayes and Wheelwright's Four-Stage Model of operations contribution sees operations as moving from implementation of strategy, through to supporting strategy, and finally driving strategy**

Stage 2: External neutrality

The first step of breaking out of Stage 1 is for the operations function to begin comparing itself with similar companies or organisations in the outside market. This may not immediately take it to the 'first division' of companies in the market, but at least it is measuring itself against its competitors' performance and trying to be 'appropriate', by adopting 'best practice' from them. Its vision is to become 'up to speed' or 'externally neutral' with similar businesses in its industry by adopting 'best practice' ideas and norms of performance from others.

Stage 3: Internally supportive

Stage 3 operations have probably reached the 'first division' in their market. They may not be better than their competitors on every aspect of operations performance but they are broadly up with the best. Yet, the vision of Stage 3 operations is to be clearly and unambiguously the very best in the market. They may try to achieve this by gaining a clear view of the company's competitive or strategic goals and developing 'appropriate' operations resources to excel in the areas in which the company needs to compete effectively. The operation is trying to be 'internally supportive' by providing a credible operations strategy.

Stage 4: Externally supportive

Stage 3 used to be taken as the limit of the operations function's contribution. Yet, the model captures the growing importance of operations management by suggesting a further stage – Stage 4. The difference between Stages 3 and 4 is subtle, but important. A Stage 4 company is one where the vision for the operations function is to provide *the* foundation for competitive success. Operations looks to the long term. It forecasts likely changes in markets and supply, and, over time, it develops the operations-based capabilities that will be required to compete in future market conditions. The operations function is becoming central to strategy-making. Stage 4 operations are creative and proactive. They are innovative and capable of adaptation as markets change. Essentially, they are trying to be 'one step ahead' of competitors in the way that they create products and services and organise their operations – what the model terms being 'externally supportive'.

The four perspectives on operations strategy

So, as an enterprise moves through each of these stages, what exactly should an operations strategy do? Different authors have slightly different views and definitions of operations strategy. However, between them, four 'perspectives' emerge:[3]

1. Operations strategy should reflect what the whole group or business wants to do in a 'top-down' manner.
2. Operations strategy should translate the enterprise's intended market position so as to provide the required objectives for operations decisions (sometimes called the 'outside-in' perspective).
3. Operations strategy should learn from day-to-day activities so as to cumulatively build strategic capabilities in a 'bottom-up' manner.
4. Operations strategy should develop its resources and processes so that their capabilities can be exploited in their chosen markets (sometimes called the 'inside-out' perspective).

None of these four perspectives alone gives the full picture of what operations strategy is. But together they provide some idea of the pressures that go to form the content of operations strategy. We will treat each of them in turn. *See* Figure 2.3.

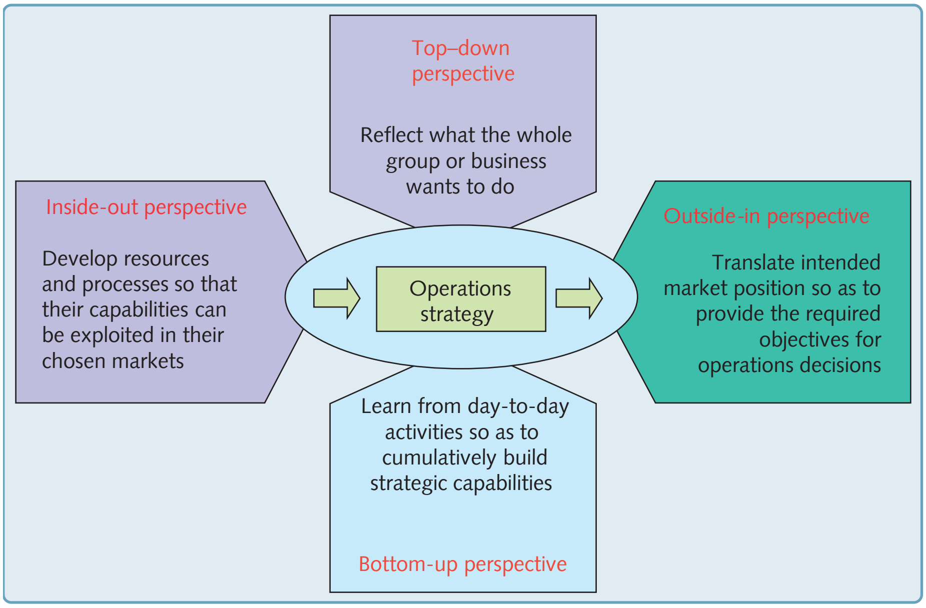

Figure 2.3 **The four perspectives on operations strategy**

DIAGNOSTIC QUESTION

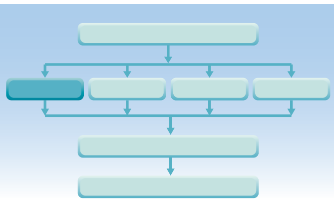

Does operations strategy reflect business strategy (top-down)?

A top-down perspective often identifies three related levels of strategy: corporate, business and functional. A corporate strategy should position the corporation in its global, economic, political and social environment. This will consist of decisions about what types of business the group wants to be in, what parts of the world it wants to operate in, how to allocate its cash between its various businesses, and so on. Each business unit within the corporate group will also need to put together its own business strategy that sets out its individual mission and objectives. This business strategy guides the business in relation to its customers, markets and competitors, and also defines its role within the corporate group of which it is a part. Similarly, within the business, functional strategies need to consider what part each function should play in contributing to the strategic objectives of the business. The operations, marketing, product/service development and other functions will all need to consider how best they should organise themselves to support the business's objectives. This is why it is often called the 'top-down' perspective on operations strategy.

OPERATIONS PRINCIPLE

Operations strategies should reflect top-down corporate and/or business objectives.

Although this rather neat relationship between the levels of corporate, business and operations strategy may seem a little 'theoretical', it is still a powerful idea. What it is saying is that in order to understand strategy at any level, one has to place it in the context of what it is trying to do (the level above) and how it is trying to do it (the level below). At any level, a good top-down perspective should provide clarity and connection. It should clarify what an operations strategy should be prioritising, and give some guidance on how the strategy is to be achieved.

Correspondence and coherence

However, developing any functional strategy from a business strategy is not a straightforward task. There are ambiguities to clarify and conflicts to be reconciled. Inevitably, business strategy consists of aggregated and approximate objectives. It should give an overall direction, but cannot spell out every detail of how a function should interpret its objectives. Yet, there should be a clear, explicit and logical connection between each functional strategy and the business strategy in which they operate. Moreover, there should also be a clear, explicit and logical connection between a functional strategy and the decisions taken within the function. In other words there should a clear *correspondence* between a business's strategy and its operations strategy, as there should also be between an operations strategy and the individual decisions taken within the operations function.

But although correspondence between the levels of strategy is necessary, it is not all that is required. Operations strategy must also be *coherent,* both with other functional strategies and within itself. Coherence means that those choices made across or within functions should not pull it in different directions. All decisions should complement and reinforce each other in the promotion of the business's and the operation's objectives. Figure 2.4 illustrates these two ideas of correspondence and coherence.[4]

EXAMPLE

Innovation at Micraytech (Part 1, top-down)

Micraytech is a metrology systems company that develops integrated measurement systems for large international clients in several industries; it is part of the Micray Group that includes several high-tech companies. It has grown through a strategy of providing products with a high degree of technical excellence and innovation, together with an ability to customise its systems and offer technical advice to its clients. The Group has set ambitious growth targets for the company over the next 5 years and has relaxed its normal 'return on sales' targets to help it achieve this. As part of this strategy, Micraytech attempted to be the first in the market with all appropriate new technical innovations. From a top-down perspective, its operations function, therefore, needed to be capable of coping with the changes that constant product innovation would bring.

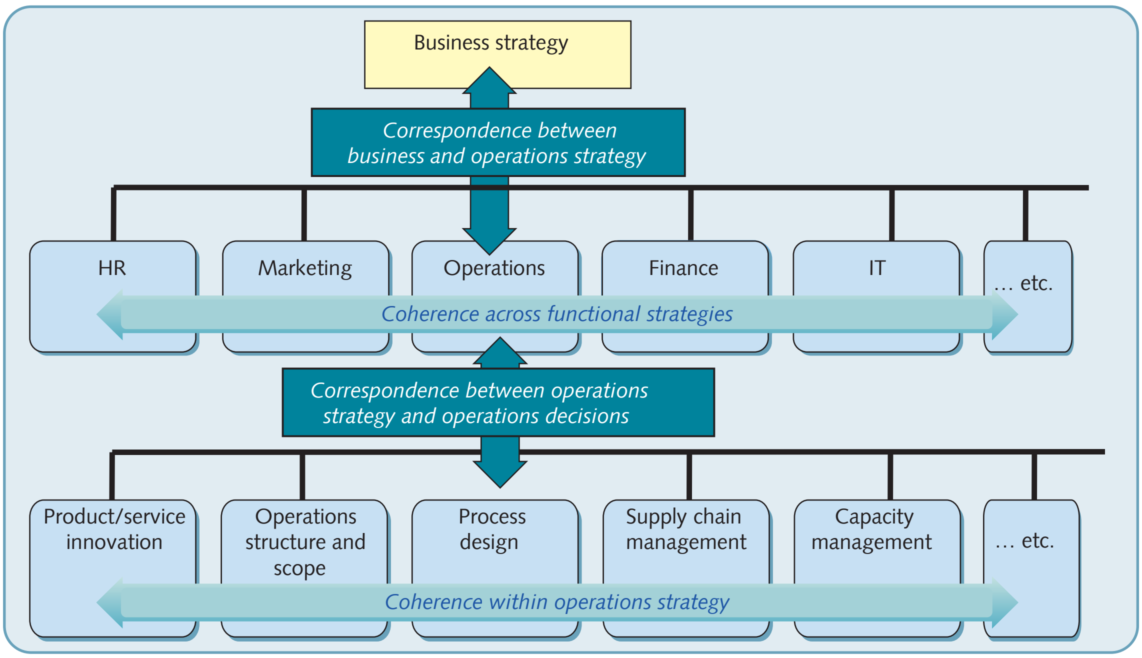

Figure 2.4 Correspondence and coherence are the two requirements of the top-down perspective of operations strategy

It developed processes that were flexible enough to develop and assemble novel components and systems, while integrating them with software innovations.

The company's operations staff realised that they needed to organise and train their staff to understand the way technology is developing so that they could put in place the necessary changes to the operation. It also needed to develop relationships with both existing and potentially new suppliers who could respond quickly when supplying new components. So the top-down logic here is that everything about the operation – its processes, staff, and its systems and procedures – must, in the short-term, do nothing to inhibit, and in the long term actively develop, the company's competitive strategy of growth through innovation.

The concepts of the 'business model' and the 'operating model'

Two concepts have emerged over the last few years that are useful in understanding the top-down perspective on operations strategy (or at least the terms are new, one could argue that the ideas are far older). These are the concepts of the 'business model' and the 'operating model'.

Put simply, a 'business model' is the plan that is implemented by a company to generate revenue and make a profit (or fulfill its social objectives if a not-for-profit enterprise). It includes the various parts and organisational functions of the business, as well as the revenues it generates and the expenses it incurs. In other words, what a company does and how they make money from doing it. More formally, it is . . . 'A conceptual tool that contains a big set of elements and their relationships and allows [the expression of] the business logic of a specific firm. It is a description of the value a company offers to one or several segments of customers and of the architecture of the firm and its network of partners for creating, marketing, and delivering this value and relationship capital, to generate profitable and sustainable revenue streams.'[5]

One synthesis of literature[6] shows that business models have a number of common elements:

1. The *value proposition* of what is offered to the market.
2. The *target customer segments* addressed by the value proposition.
3. The communication and *distribution channels* to reach customers and offer the value proposition.
4. The *relationships* established with customers.
5. The *core capabilities* needed to make the business model possible.
6. The *configuration of activities* to implement the business model.
7. The *partners* and their motivations for coming together to make a business model happen.
8. The *revenue streams* generated by the business model constituting the revenue model.
9. The *cost structure* resulting from the business model.

One can see that this idea of the business model is broadly analogous to the idea of a 'business strategy', but implies more of an emphasis on *how* to achieve an intended strategy as well as exactly *what* that strategy should be.

The concept of an 'operating model' is more operational than that of a 'business model' and there is no universally agreed definition. Here we define it as a . . . 'High-level design of the organisation that defines the structure and style which enables it to meet its business objectives'.[7] Ideally, an operating model should provide a clear, 'big picture' description of what the organisation does and how it does it. It defines how the critical work of an organisation is carried out. It should provide a way to examine the business in terms of the key relationships between business functions, processes and structures that are required for the organisation to fulfil its mission. Unlike the concept of a business model, which often assumes a profit motive, the operating model philosophy can be applied to organisations of all types – including large corporations, not-for-profit organisations and the public sector.

Again, there is no universally agreed list of elements that an operating model should include and different organisations focus on different things, but many of the following elements are often included:

- Key performance indicators (KPIs) – with an indication of the relative importance of performance objectives.
- Core financial structure – P&L, new investments and cash flow.
- The nature of accountabilities for products, geographies, assets. etc.
- The structure of the organisation – sometimes expressed as capability areas rather than functional roles.
- Systems and technologies.
- Processes, responsibilities and interactions.
- Key knowledge and competence.

Note two important characteristics of an operating model. First, it does not respect conventional functional boundaries as such. In some ways the concept of the operating model reflects the idea that we proposed in Chapter 1. Namely, that all managers are operations managers and all functions can be considered as operations because they comprise processes that deliver some kind of service. An operating model is like an operations strategy, but applied across all functions and domains of the organisation. Second, there are clear overlaps between the 'business model' and the 'operating model'. The main difference being that an operating model focuses more on how an overall business strategy is to be achieved. Also, operating models are rarely designed from first principles. Some kind of understood 'way of doing things' will already exist. This is why operating models often have an element of implied change or transformation of the organisation's resources and processes. Often the term 'target operating model' (TOM) is used to describe the way the organisation should operate in the future if it is going to achieve its objectives and make a success of its business model. Figure 2.5 illustrates the relationship between business and operating models.

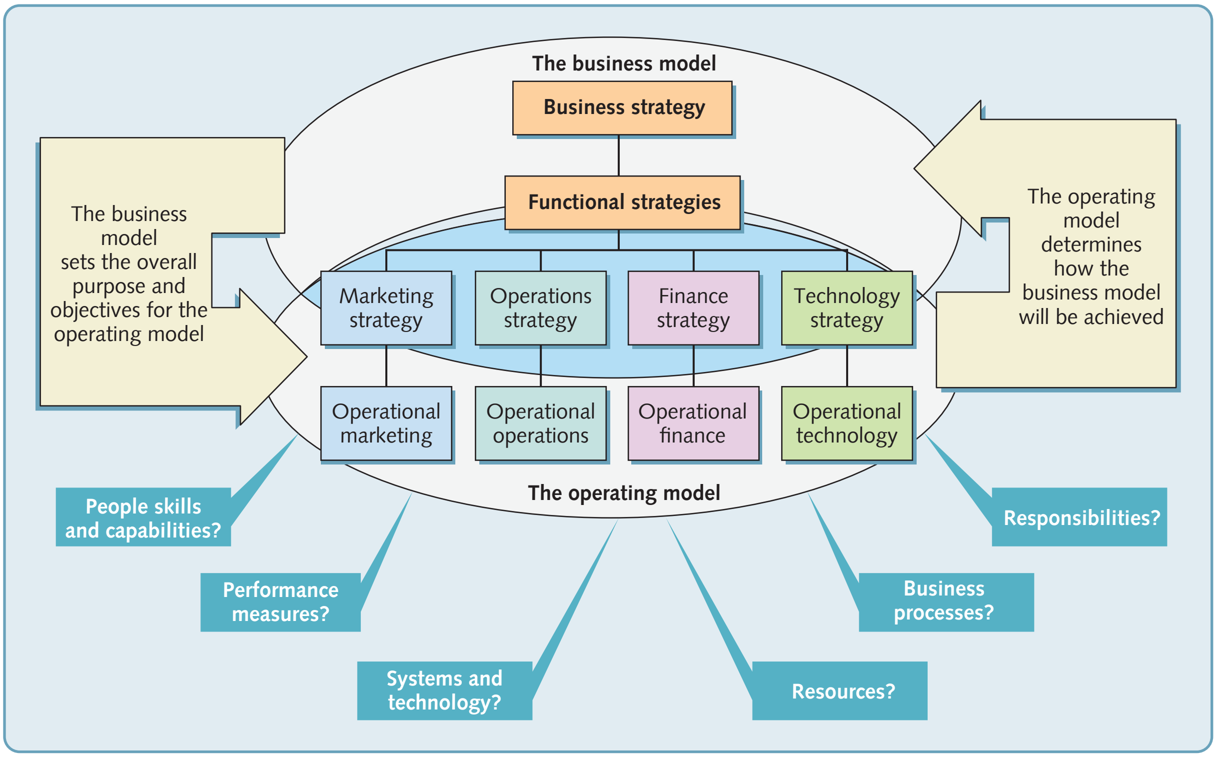

Figure 2.5 **The concepts of 'the business model' and the 'operating model' overlap – with the operating model indicating how processes, resources, technology, people, measures and responsibilities are to be organised to support the business model**

DIAGNOSTIC QUESTION

Does operations strategy align with market requirements (outside-in)?

Any operations strategy should reflect the intended market position of the business. Companies compete in different ways; some compete primarily on cost, others on the excellence of their products or services, others on high levels of customer service, and so on. The operations function should therefore respond to this by providing the ability to perform in a manner that is appropriate for the intended market position. This is a market (or outside-in) perspective on operations strategy.

An operations strategy should define operations performance objectives

Operations add value for customers and contribute to competitiveness by being able to satisfy the requirements of its customers. The most useful way to do this is to use the 'operational' performance measures that we briefly discussed in the previous chapter. All of these, to a greater or lesser extent, will affect customer satisfaction and define market positioning in terms that have meaning in an operations context. As a reminder, the five performance objectives are:

- **Quality** – producing error-free goods and services that are 'fit for their purpose'.
- **Speed** – minimising the time between a customer asking for goods and services and the customer receiving them in full.
- **Dependability** – keeping the delivery promises that have been made to customers.
- **Flexibility** – the ability to vary or adapt the operation's activities to cope with unexpected circumstances or to give customers individual treatment, or to introduce new products or services.
- **Cost** – producing goods and services at a cost that enables them to be priced appropriately for the market, while still allowing a return to the organisation (or, in a not-for-profit organisation, that give good value to the tax-payers or whoever is funding the operation).

The exact meaning of performance objectives is different in different operations

Different operations will have different views of what each of the performance objectives actually mean. Table 2.2 looks at how two operations, an insurance company and a steel plant, define each performance objective. For example, the insurance company sees quality as being at least as much about the manner in which their customers relate to their service, as it does about the absence of technical errors. The steel plant, on the other hand, while not ignoring quality of service, primarily emphasises product-related technical issues. Although, they are selecting from the same pool of factors which together constitute the generic performance objective, they will emphasise different elements.

OPERATIONS PRINCIPLE

Operations performance objectives can be grouped together as quality, speed, dependability, flexibility and cost.

Sometimes operations may choose to re-bundle elements using slightly different headings. For example, it is not uncommon in some service operations to refer to 'quality of service' as representing all the competitive factors we have listed under quality *and* speed *and* dependability (and sometimes aspects of flexibility). For example, information

Table 2.2 Aspects of performance objectives for two operations

Insurance company *Aspects of each performance objectives include. . .*	*Performance objectives*	*Steel plant* *Aspects of each performance objectives include. . .*
• Professionalism of staff • Friendliness of staff • Accuracy of information • Ability to change details in future	**Quality**	• Percentage of products conforming to their specification • Absolute specification of products • Usefulness of technical advice
• Time for call centre to respond • Prompt advice response • Fast quotation decisions • Fast response to claims	**Speed**	• Lead time from enquiry to quotation • Lead time from order to delivery • Lead time for technical advice
• Reliability of original promise date • Customers kept informed	**Dependability**	• Percentage of deliveries 'on-time, in-full' • Customers kept informed of delivery dates
• Customisation of terms of insurance cover • Ability to cope with changes in circumstances, such as level of demand • Ability to handle wide variety of risks	**Flexibility**	• Range of sizes, gauges, coatings, etc. possible • Rate of new product introduction • Ability to change quantity, composition and timing of an order
• Premium charged • Arrangement charges • 'No-claims' deals • 'Excess' charges	**Cost**	• Price of products • Price of technical advice • Discounts available • Payment terms

network operations use the term 'Quality of Service' (QoS) to describe their goal of providing guarantees on the ability of a network to deliver predictable results. This is often specified as including uptime (dependability), bandwidth provision (dependability and flexibility), latency or delay (speed of throughput) and error rate (quality). In practice, the issue is not so much one of universal definition but rather consistency within one, or a group of operations. At the very least it is important that individual companies have it clear in their own minds how each performance objective is to be defined.

OPERATIONS PRINCIPLE

The interpretation of the five performance objectives will differ between different operations.

The relative priority of performance objectives differs between businesses

The idea behind the outside-in perspective is that businesses that compete in different ways should want different things from their operations functions. Therefore, not every operation will apply the same priorities to its performance objectives. So there should be a clear logical connection between the competitive stance of a business and its operations objectives. For example, a business that competes primarily on low prices and 'value for money' should place emphasis on operations objectives such as cost, productivity and efficiency; one that competes on a high degree of customisation of its services or products should place an emphasis on flexibility, and so on. Many successful companies understand the importance of making this connection between their message to customers and the operations performance objectives that they emphasise. For example,[8]

> *'We focus on providing the best user experience possible . . . through continued iteration on difficult problems, we've been able to solve complex issues and provide continuous improvements to a service . . . fast is better than slow . . . through innovation and iteration, we aim to take things that work well and improve upon them . . . we see being great at something as a starting point, not an end point'*
>
> (Google)

'Our management principle is the commitment to quality and reliability . . . to deliver safe and innovative products and services . . . and to improve the quality and reliability of our businesses.'

(Komatsu)

OPERATIONS PRINCIPLE

The relative importance of the five performance objectives depends on how the business competes in its market.

'A level of quality durability and value that's truly superior in the market place. . . the principle that what is best for the customer is also best for the company. . . (our) . . . customers have learnt to expect a high level of service at all times – from initiating the order, to receiving help and advice, to speedy shipping and further follow-up where necessary. . . (our) . . . employees 'go that extra mile'.'

(Lands' End)

Order winners and qualifiers

OPERATIONS PRINCIPLE

Operations strategy should reflect the requirements of the business's markets.

A particularly useful way of determining the relative importance of competitive factors is to distinguish between what have been termed 'order-winners' and 'qualifiers'.[9] Figure 2.6 shows the difference between order-winning and qualifying objectives in terms of their utility, or worth, to the competitiveness of the organisation. The curves illustrate the relative amount of competitiveness (or attractiveness to customers) as the operation's performance varies.

Order-winners are those things that directly and significantly contribute to winning business. They are regarded by customers as key reasons for purchasing the product or service. Raising performance in an order-winner will either result in more business or improve the chances of gaining more business. Order-winners show a steady and significant increase in their contribution to competitiveness as the operation gets better at providing them.

Qualifiers may not be the major competitive determinants of success, but are important in another way. They are those aspects of competitiveness where the operation's performance has to be above a particular level just to be considered by the customer. Performance below this 'qualifying' level of performance may disqualify the operation from being considered by customers. But any further improvement above the qualifying level is unlikely to gain the company much competitive benefit. Qualifiers are those things that are generally expected by customers. Being great at them is unlikely to excite them, but being bad at them can disadvantage the competitive position of the operation.

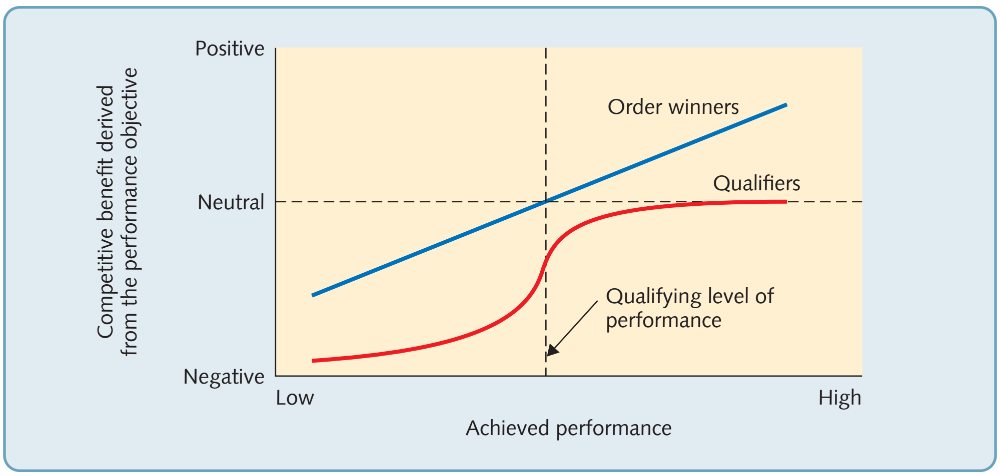

Figure 2.6 Order winners and qualifiers. Order winners gain more business the better you are. Qualifiers are the 'givens' of doing business

Different customer needs imply different objectives

OPERATIONS PRINCIPLE
Different customer needs imply different priorities of performance objectives.

If, as is likely, an operation produces goods or services for more than one customer group, it will need to determine the order-winners and qualifiers for each group. For example, Table 2.3 shows two 'product' groups in the banking industry. Here the distinction is drawn between the customers who are looking for banking services for their private and domestic needs and the corporate customers who need banking services for their (often large) businesses.

The product/service life cycle influence on performance objectives

One way of generalising the market requirements that operations need to fulfil is to link it to the life cycle of the products or services that the operation is producing. The exact form of product/service life cycles will vary, but generally they are shown as the sales volume passing through four stages – introduction, growth, maturity and decline. The important implication of this for operations management is that products and services will require operations strategies in each stage of their life cycle (see Figure 2.7).

- **Introduction stage** – When a product or service is first introduced, it is likely to be offering something new in terms of its design or performance. Given the market uncertainty, the operations management of the company needs to develop the flexibility to cope with these changes and the quality to maintain product/service performance.
- **Growth stage** – In the growing market, standardised designs emerge that allow the operation to supply the rapidly growing market. Keeping up with demand through rapid and dependable response and maintaining quality levels will help to keep market share as competition starts to increase.

Table 2.3 Different banking services require different performance objectives

	Retail banking	*Corporate banking*
Products	Personal financial services such as loans and credit cards	Special services for corporate customers
Customers	Individuals	Businesses
Product range	Medium but standardised, little need for special terms	Very wide range, many need to be customised
Design changes	Occasional	Continual
Delivery	Fast decisions	Dependable service
Quality	Means error-free transactions	Means close relationships
Volume per service type	Most service are high volume	Most services are low volume
Profit margins	Most are low to medium, some high	Medium to high
Order winners	Price Accessibility Speed	Customisation Quality of service Reliability
Qualifiers	Quality Range	Speed Price
Performance objectives emphasised within the processes that produce each service	Cost Speed Quality	Flexibility Quality Dependability

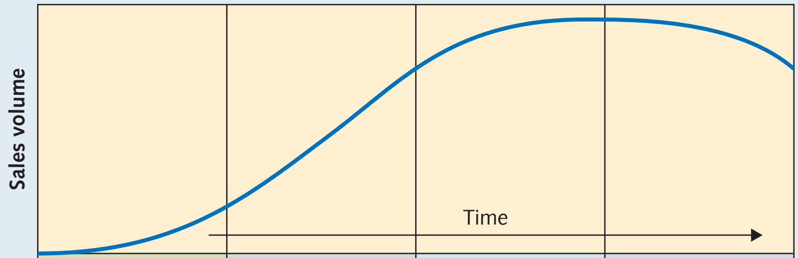

	Introduction	Growth	Maturity	Decline
	Product/service first introduced to market	Product/service gains market acceptance	Market needs start to be fulfilled	Market needs largely met
Volume	Slow growth in sales	Rapid growth in sales volumes	Sales slow down and level off	Sales decline
Customers	Innovators	Early adopters	Bulk of market	Laggards
Competitors	Few/none	Increasing numbers	Stable number	Declining numbers
Variety of product/service design	Possible high customisation or frequent design changes	Increasingly standardised	Emerging dominant types	Possible move to commodity standardisation
Likely order-winners	Product/service characteristics, performance or novelty	Availability of quality products/services	Low price Dependable supply	Low price
Likely qualifiers	Quality Range	Price Range	Range Quality	Dependable supply
Dominant process performance objectives	Flexibility Quality	Speed Dependability Quality	Cost Dependability	Cost

Figure 2.7 The effects of the product/service life cycle on the operation and its process performance objectives

- **Maturity stage** – Eventually demand starts to level off as the market becomes dominated by a few larger companies with standardised designs. Competition will probably emphasise price or value for money. So operations will be expected to get the costs down in order to maintain profits or to allow price cutting, or both. So, cost and productivity issues, together with dependable supply, are likely to be the operation's main concerns.
- **Decline stage** – After time, sales will decline. To the companies left there might be a residual market, but if capacity in the industry lags demand, the market will be dominated by price competition; therefore cost cutting continues to be important.

EXAMPLE

Innovation at Micraytech (Part 2, outside-in)

The Micray Group sees a major growth opportunity for Micraytech by continually incorporating technological innovations in its product offerings. However, Micraytech's marketing management know that this can be achieved by focusing on one or both of two distinct markets. The first is the market for 'individual metrology devices'. These are 'stand-alone' pieces of equipment bought by all types of industrial customers. It had traditionally been the company's main market. The second market was that for 'integrated metrology systems'. These were larger, more complex, more expensive (and higher margin) offerings that were customised to individual customers' requirements. The two types of offering had overlapping, but different characteristics. 'Individual metrology devices' competed on their technical performance and reliability, together with relatively short delivery times compared with competitors. The 'integrated metrology systems' offerings currently accounted for only a small part of the company's sales, but it was a market that was forecast to grow substantially. The customers for these systems were larger manufacturers who were investing in more automated technologies and required metrology systems that could be integrated into their processes. From an 'outside-in' perspective, if it was to take advantage of this emerging market, Micraytech would have to learn how to work more closely with both their direct customers and the firms that were supplying their customers with the automated technologies. In addition to Micraytech's traditional technical skills, it would have to increase its software development, data exchange and client liaison, skills.

DIAGNOSTIC QUESTION

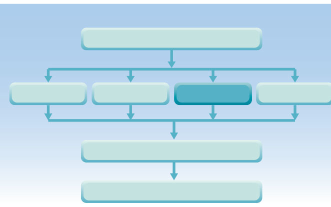

Does operations strategy learn from operational experience (bottom-up)?

Although it is a convenient way of thinking about strategy, the top-down hierarchical model does not represent the way strategies are always formulated in practice. When any group is reviewing its corporate strategy, it will also take into account the circumstances, experiences and capabilities of the various businesses that form the group. This is sensible. Operations strategy should always reflect operational reality. So, businesses, when reviewing their strategies, should consult the individual functions within the business about their constraints and capabilities. They may also incorporate the ideas that come from each function's day-to-day experience. In fact, many strategic ideas emerge over time from operational experience rather than originating exclusively at a senior level.

Sometimes companies move in a particular strategic direction because the ongoing experience of providing products and services to customers at an operational level convinces them that it is the right thing to do. There may be no formal high-level decision-making that examines alternative strategic options and the one that provides the best way forward is chosen. Instead, a general consensus emerges from the operational experience. The 'high level' strategic decision-making, if it occurs at all, may simply confirm the consensus and provide the resources to make it happen effectively. This is sometimes called the concept of 'emergent strategies'.[10] It sees strategy-making, at least partly, as a relatively unstructured and fragmented process to reflect the fact that the future is at least partially unknown and unpredictable.

OPERATIONS PRINCIPLE

Operations strategy should reflect bottom-up experience of operational reality.

This perspective on operations strategy reflects how things often happen, but at first glance it seems less useful in providing a guide for specific decision-making. Yet, while emergent strategies are less easy to categorise, the principle governing a bottom-up perspective is clear: operation's objectives and action, should be shaped at least partly, by the knowledge it gains from its day-to-day activities. Certainly, encouraging strategic ideas to emerge from the operational experience is not an abandonment of strategic responsibility by senior management. Rather it is an acceptance that worthwhile ideas can come from those who work at the operational level of the business. In fact, not accepting it would be to waste one of the most insightful sources of practical understanding of operational reality. From a bottom-up perspective, the key virtues required for shaping strategy from the bottom up is an ability to learn from experience. More specifically, the bottom-up perspective should involve:

- Capturing the learning that should come from routine, operations activities, and
- Transforming that learning into strategically valuable knowledge.

Another advantage of paying attention to the bottom-up perspective is that it is at the operational day-to-day level experience where trends often first become evident. And, for businesses operating in unstable or unpredictable environment, this can be particularly important. The other three perspectives of operations strategy can take time to detect trends in how markets are moving. The bottom-up element is more 'plugged in' to everyday experience.

EXAMPLE

Innovation at Micraytech (Part 3, bottom-up)

Over time, as its operations strategy developed, Micraytech discovered that continual product and system innovation was having the effect of dramatically increasing its costs. And, although it was not competing on low prices, and nor was it under pressure from The Group to achieve high rates of return on sales, its rising costs were impacting profitability to an unacceptable degree. Also there was some evidence that continual updating of product and system specifications were confusing some customers. Partially in response to customer requests, the company's system designers started to work out a way of 'modularising' their system and product designs. This allowed one part of the system to be updated for those customers who valued the functionality that the innovation could bring, without interfering with the overall design of the main body of the system of which the module was a part. Over time, this approach became standard design practice within the company. Customers appreciated the extra customisation, and modularisation reduced operations costs. Note that this strategy emerged from the company's experience. It was a pure bottom-up approach. Initially, no top-level board decision was taken to initiate this practice. Nevertheless, it emerged as the way in which the company's design engineers learned from their experience and used that learning to build their knowledge of how to lower some of the costs of innovation.

DIAGNOSTIC QUESTION

Does operations strategy develop the capability of its resources and processes (inside-out)?

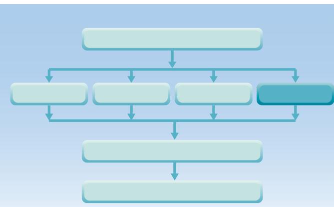

Operations strategy must do more than simply meet the short-term needs of the market (important though this is). The processes and resources within operations also need to be developed in the long term to provide the business with a set of competencies or capabilities (we use the two words interchangeably). A capability in this context is the 'know how' that is embedded within the business's resources and processes. These capabilities may be built up over time as the result of the experiences of the operation (bottom-up), or they may be bought-in or acquired. If they are refined and integrated they can form the basis of the business's ability to offer unique and/or 'difficult to imitate' products and services to its customers. This idea of the basis of long-term competitive capabilities deriving from the operation's resources and processes is called the resource, or inside-out, perspective on operations strategy.[11]

Operations strategy should build operations capabilities

Building operations capabilities means understanding the existing resources and processes within the operation, starting with the simple questions, what do we have, and what can we do? However, trying to understand an operation by listing its resources alone is like trying to understand an automobile by listing its component parts. To understand an automobile we need to describe how the component parts form its internal mechanisms. Within the operation, the equivalents of these mechanisms are its *processes.* Yet, even a technical explanation of an automobile's mechanisms does not convey its style or 'personality'. Something more is needed to describe these. In the same way, an operation is not just the sum of its processes. It also has *intangible* resources. An operation's intangible resources include such things as:

> **OPERATIONS PRINCIPLE**
> ***The long-term objective of operation strategy is to build 'operations-based capabilities'.***

- its relationship with suppliers and the reputation it has with its customers
- its knowledge of and experience in handling its process technologies
- the way its staff can work together in new product and service development
- the way it integrates all its processes into a mutually supporting whole.

These intangible resources may not be as evident within an operation, but they are important and often have real value. And both tangible and intangible resources and processes shape its capabilities. The central issue for operations management, therefore, is to ensure that its pattern of strategic decisions really does develop appropriate capabilities.

An operations strategy should identify the broad decisions that will help the operation build its capabilities

Few businesses have the resources to pursue every single action that might improve their operations performance. So an operations strategy should indicate broadly how the operation might best achieve its performance objectives. For example, a business might specify that it will attempt to reduce its costs by aggressive outsourcing of its non-core business processes and by investing in more efficient technology. Or, it may declare that it intends to offer a more customised set of products or services through adopting a modular approach to its product or service design. The balance here is between a strategy that is overly restrictive in specifying how performance objectives are to be achieved, and one that is so open that it gives little guidance as to what ideas should be pursued.

There are several categorisations of operations strategy decisions.[12] Any of them are valid if they capture the key decisions. Table 2.4 illustrates some of the broad operations strategy decisions that fall within each category.

The resource-based view

The idea that building operations capabilities is a particularly important objective of operations strategy is closely linked with the 'resource based view' (RBV) of the firm. This holds that businesses with an 'above-average' strategic performance are likely to have gained their sustainable competitive advantage because of their core competences (or capabilities). This means that the way an organisation inherits, or acquires, or develops its operations resources will, over the long term, have a significant impact on its strategic success. The RBV differs in its approach from the more traditional view of strategy that sees companies as seeking to protect their competitive advantage through their control of the market. For example, they may do this by creating *barriers to entry* through product differentiation, or making it difficult for customers to switch to competitors, or controlling the access to distribution channels (a major barrier to entry in

Table 2.4 Some strategic decisions that may be addressed in an operations strategy

• How should the operation decide which products or services to develop and how to manage the development process?
• Should the operation outsource some of its activities, or take more activities in-house?
• Should the operation expand by acquiring its suppliers or its customers? If so, which ones should it acquire?
• How many geographically separate sites should the operation have?
• Where should operations sites be located?
• What activities and capacity should be allocated to each site?
• What broad types of technology should the operation be using?
• How should the operation be developing its people?
• What role should the people who staff the operation play in its management?
• How should the operation forecast and monitor the demand for its offerings?
• How should the operation adjust its activity levels in response to demand fluctuations?
• How should the operation monitor and develop its relationship with its suppliers?
• How much inventory should the operation have and where should it be located?
• What system should the operation use to coordinate its activities?
• How should the operation's performance be measured and reported?
• How should the operation ensure that its performance is reflected in its improvement priorities?
• Who should be involved in the improvement process?
• Should improvement be continuous, or radical, or both?
• How should the improvement process be managed?
• How should the operation maintain its resources so as to prevent failure?
• How should the operation ensure continuity if a failure occurs?

gasoline retailing, for example, where oil companies own their own retail stations). By contrast, the RBV sees firms being able to protect their competitive advantage through *barriers to imitation,* that is, by building up 'difficult-to-imitate' resources. Some of these 'difficult-to-imitate' resources are particularly important, and can be classified as 'strategic' if they exhibit the following properties:[13]

- *They are scarce* – Scarce resources, such as specialised production facilities, experienced engineers, proprietary software, etc. can underpin competitive advantage.
- They are imperfectly mobile – Some resources are difficult to move out of a firm. For example, resources that were developed in-house, or are based on the experience of the company's staff, or are interconnected with the other resources in the firm cannot be traded easily.
- *They are imperfectly imitable and imperfectly substitutable* – it is not enough only to have resources that are unique and immobile. If a competitor can copy these resources or, replace them with alternative resources, then their value will quickly deteriorate. The more the resources are connected with process knowledge embedded deep within the firm, the more difficult they are for competitors to understand and to copy.

The VRIO framework

The most common (and useful) way of applying the RBV has become known as the VRIO framework.[14] It was first developed by Barney in the 1990s[15] (who originally identified the idea of resources needing to be scarce, imperfectly mobile, imperfectly imitable and imperfectly substitutable) but later modified to make it more useful for practitioners. In this framework, the four questions to ask about any potentially strategic resource are:

1. ***Is the resource valuable?*** – Is it possible to identify specific and definable competitive value from the resources?
2. ***Is the resource rare?*** – Do you have (or have access to) resources that your competitors do not? Some theorists define the idea of 'rarity' as when a business has a resource that is unequivocally unique, but for all practical purposes, a resource is 'rare' if it is, at least, in short supply and likely to remain so.

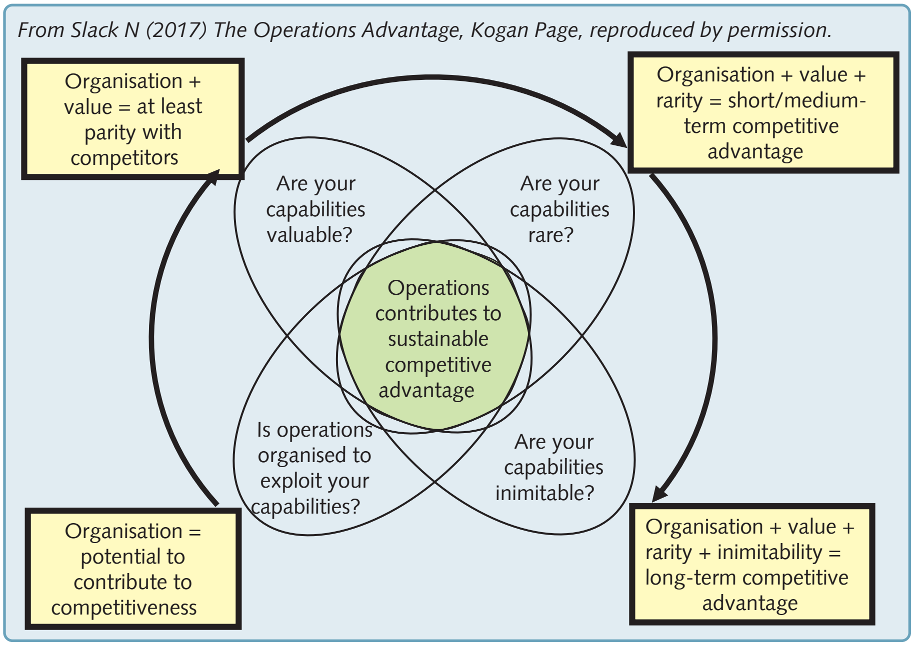

Figure 2.8 The four features of Barney's VRIO framework

Source: From Slack, N. (2017) *The Operations Advantage*, Kogan Page, reproduced by permission.

3. ***Is the resource costly to imitate?*** – Do you have resources that competitors cannot imitate, purchase or find a suitable alternative to, at a realistic cost or in a realistic time frame? 'Imitability' can be either because competitors can copy your resources and processes directly, or because they can find an acceptable substitute for them.
4. ***Is the firm organised to capture the value of the resource?*** – Do you have within your business the systems, culture, capacity and motivation to exploit any capabilities embedded in your resources and processes? A firm must have the formal reporting and control mechanisms, leadership, and the informal and cultural environment that allows the strategic resources to develop.

There are two important points to remember about the VRIO framework. First, all these factors are time dependent. A capability may be currently valuable now, but competitors are unlikely to stand still. In addition, rarity and inimitability are not absolutes and, with time, can be undermined by competitor activity. Even the ability to exploit capabilities can erode if operations leadership is lacking. Second, although the conventional order in which to treat each of these elements is as we have done here (which is why it's called the VRIO framework), it maybe is best to think of the 'O' of 'organisation' to be a necessary prerequisite. Without the ability to exploit strategic resources, they are of little use. However, with effective organisation there is the potential for operations resources to contribute to competitiveness. If their capabilities are also valuable, then parity with competitors should be possible. With the addition of rarity, a short- to medium-term competitive advantage is possible. With the addition of inimitability, competitors will find it difficult to match capabilities in anything but the long term. This sequence is shown in Figure 2.8.

EXAMPLE

Innovation at Micraytech (Part 4, inside-out)

The modular approach to product design proved to be a big success for Micraytech. However, it posed two challenges for the company's operations. First, the technical aspects of integrating some of the more sophisticated modules proved difficult. This only affected a small proportion of customers, but they were the ones that were willing to pay premium prices for their systems. The only potential solution was to attempt to develop the interface modules that would allow previously incompatible modules to be integrated. When this solution was first proposed the relevant skills were not present in the company. It had to recruit specialist engineers to start on the design of the interfaces. During this design process the company realised that it could potentially open up a new market. As the firm's Chief Operating Officer (COO) put it, '*If we designed the interfaces carefully, we could not only integrate all of our own in-house modules, we could also integrate other firms' instruments into our systems*'. This led to the second set of challenges; to develop relationships with possible suppliers, who might very well be competitors in some markets, so that they were willing to supply their equipment for inclusion in Micraytech's systems. Not only this, but also the firm had to ensure that the internal processes, of its sales engineers consulting with clients, its design department designing the system to clients' needs, and its procurement managers negotiating with equipment suppliers, all operated seamlessly. '*The success that we have enjoyed can be put down to two key capabilities. The first was to buy in the engineering skills to create technically difficult interfaces. That led to us understanding the value that could be gained from a seamless internal and external supply chain. Both of these capabilities are not totally impossible for other firms to copy, but they would be very difficult for them to get to our level of excellence*.' (COO Micraytech)

DIAGNOSTIC QUESTION

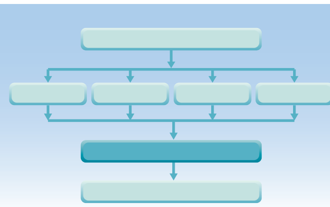

Are the four perspectives of operations strategy reconciled?

As we stressed earlier, none of the four perspectives alone can give a full picture of any organisation's operations strategy. But together they do provide a good idea of how its operations are contributing strategically. For example Figure 2.9 brings together the four perspectives of the Micraytech operations strategy. For Micraytech, the four perspectives seem to be reasonably compatible, with its operations strategy fitting together from whichever perspective is chosen. In other words, each perspective is 'reconciled' with the others. This is one of the conditions for an effective operations strategy – the four perspectives must be reconciled.

The operations strategy matrix

The operations strategy matrix is one method of checking the reconciliation between the inside-out and outside-in perspectives. It brings together (a) market requirements and (b) operations resources to form the two dimensions of a matrix. It describes operations strategy as the intersection of a company's performance objectives and the strategic decisions that it makes. In fact, there are several intersections between each performance objective and each decision area (however one wishes to define them). If a business thinks that it has an operations strategy, then it should have a coherent explanation for each of the cells in the matrix. That is, it should be able to explain and reconcile the intended links between each performance objective and each decision area. The process of reconciliation takes place between *what* is required from the operations function (performance objectives), and *how* the operation tries to achieve this through the set of choices made (and the capabilities that have been developed) in each decision area.

OPERATIONS PRINCIPLE

An operation's strategy should articulate the relationship between operations objectives and the means of achieving them.

Figure 2.10 shows a simplified example of how the operations strategy matrix can be used. A parcel courier service competes primarily on its quality and dependability of service, with price (cost) and innovation also being fairly important. The range of services offered is not unimportant, but not of prime concern. Figure 2.10 illustrates that the company believes its

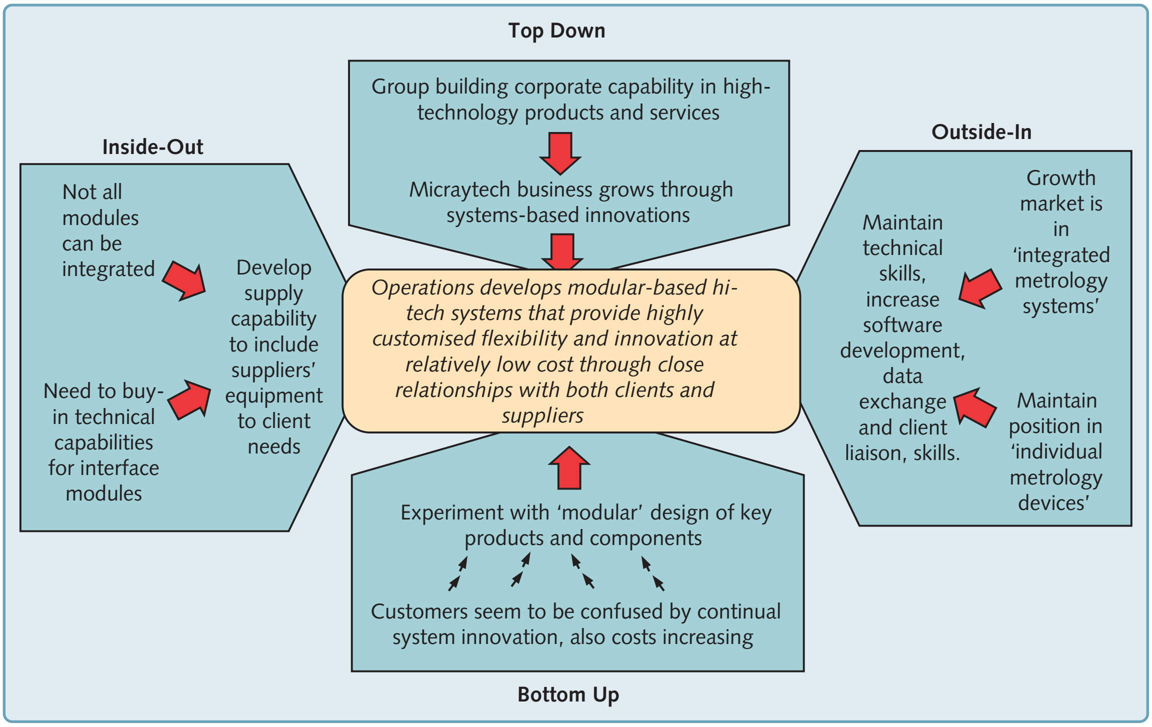

Figure 2.9 Top-down, outside-in, bottom-up and inside-out perspectives of the Micraytech operations strategy

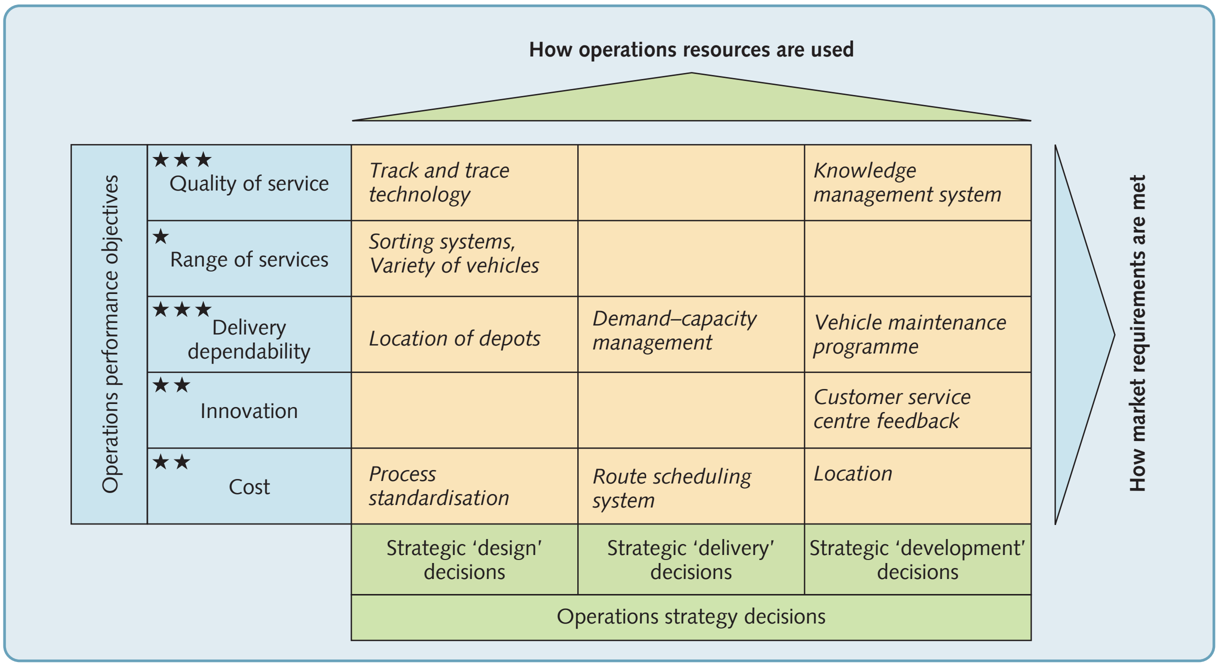

Figure 2.10 The operations strategy matrix defines operations strategy by the intersections of performance objectives and operations decisions – in this case for a parcel delivery courier

quality of service is going to be influenced largely by investment in 'track and trace' technology (that allows customers to check where their delivery is) and its knowledge management system (that allows improvements in its processes to be recorded and shared). Other key intersections are as illustrated in the figure. Note that not all cells are occupied. This is not because there is no relationship between the performance objectives and the decisions that these cells represent, it is that the decisions are not seen as being particularly important in the context of the whole strategy.

Reconciling market requirements and operations capabilities over time – the line of fit model

The operations strategy matrix is a good model for testing whether market requirements and the operations capability perspectives fit together. It makes explicit the specific aspects of market requirements (quality, speed, dependability, flexibility, cost, etc.) and the decisions that support operations capability (design, delivery and development). The disadvantage is that it gives little sense of the dynamics of reconciliation – how the balance between market requirements and the operations capability changes over time. This is where the 'line of fit' model is useful. It is based on the idea that ideally, there should be a reasonable degree of alignment, or 'fit' between the requirements of the market and the capabilities of the operation. Figure 2.11 illustrates the concept of fit diagrammatically. The vertical dimension represents the (outside-in) nature of market requirements either because they reflect the intrinsic needs of customers, or because their expectations have been shaped by the firm's marketing activity. This includes such factors as, the strength of the brand or reputation of the degree of differentiation, or the extent of market promises. Movement along the dimension indicates a broadly enhanced level of market 'performance'. The horizontal scale represents the (inside-out) nature of the firm's operations resource and process capability. This includes such things as the performance of the operation in terms of its ability to achieve competitive objectives, the efficiency with which it uses its resources and the ability of the firm's resources to underpin its business processes. Movement along the dimension broadly indicates an enhanced level of 'operations capability'.

If market requirements and operations capability of an operation are aligned it would diagrammatically be positioned on the 'line of fit' in Figure 2.11. 'Fit' is to achieve an approximate balance between 'market requirements' and 'operations capability'. So when fit is achieved, firms' customers do not need, or expect, levels of operations capability that cannot be supplied. Nor does the operation have strengths that are either inappropriate for market needs or remain unexploited in the market.

OPERATIONS PRINCIPLE

An operation positioned on the 'line of fit' has operations capabilities that match its market requirements.

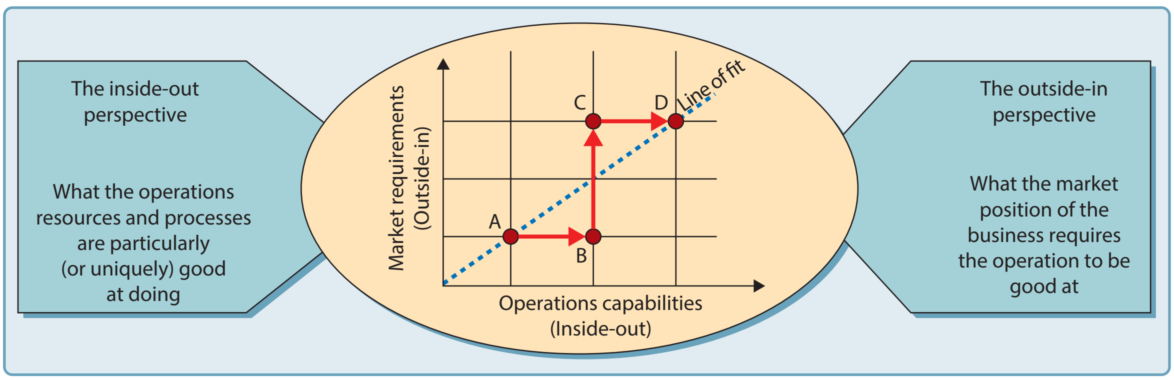

Figure 2.11 **The 'line of fit' model shows how operations strategy attempts to reconcile market requirements and operations capabilities over time**

A company that has position A in Figure 2.11 has achieved 'fit' in so much as its operations capabilities are aligned with its market requirements, yet both are at a relatively low level. In other words, the market does not want much from the business, which is just as well because its operation is not capable of achieving much. Over time its ambition is to move to position D, where it has also achieved 'fit', but at a much higher level. Other things being equal, this will be a more profitable position that position A.

However, like most strategic improvement, the company cannot always guarantee to keep on the 'line of fit' as it moves from A to D over time. In this case, it first improves its operations capability without exploiting its enhanced capability in its market (position B). This could be seen as a 'waste' of its potential to adopt a more ambitious (and possibly profitable) market position. Realising this, the company revises its marketing strategy to promote itself as being able to maintain a much higher level of market performance (position C). Unfortunately, these new promises to its market are not matched by its operations capabilities. The company is again away from the line of fit. In fact, position C is possibly even more damaging than position B. The risk now is that the company's market reputation will erode until it can improve its operations capabilities to bring it back to the line of fit (position D).

The issue that is highlighted by positioning operations strategy relative to the line of fit is that progress cannot always be a smooth trajectory that achieves perfect balance between market requirements and operations capability. Furthermore, when an operation deviates from the line of fit there are predictable consequences. A position below the line of fit means that the operation is failing to exploit its operations capabilities. A position above the line of fit means that it risks damaging its reputation or brand by failing to live up to its market promises.

EXAMPLE

Nokia, a failure to change[16]

Only a decade ago, Nokia was the king of the mobile phone business – and it was a good business to be in, with double-digit growth year on year. Nokia was omnipresent and omnipowerful, a pioneer that had supplied the first mass wave of the expanding mobile phone industry. They dominated the market in many parts of the world and the easily recognisable Nokia ring-tone echoed everywhere from boardrooms to shopping malls. So why did this, once dominant, company eventually sink to the point where it was forced to sell its mobile communications business to Microsoft in 2013? The former Nokia CEO, Jormal Ollila, admitted that Nokia made several mistakes, but the exact nature of those mistakes is a point of debate amongst business commentators. Julian Birkinshaw, a Professor at London Business School dismisses some of the most commonly cited reasons. Did they lose touch with their customers? Well, yes, but by definition that must hold for any company whose sales drop so drastically in the face of thriving competitors. Did they fail to develop the necessary technologies? No. Nokia had a prototype touchscreen before the iPhone was launched, and its smartphones were technologically superior to anything Apple, Samsung, or Google had to offer for many years. Did they not recognise that the basis of competition was shifting from the hardware to the ecosystem? (A technology ecosystem in this case is a term used to describe the complex system of interdependent components that work together to enable mobile technology to operate successfully.) Not really. The 'ecosystem' battle began in the early 2000s, with Nokia joining forces with Ericsson, Motorola and Psion to create Symbian as a platform technology that would keep Microsoft at bay.

Where they struggled was in relying on an operations strategy that failed to allocate resources appropriately and could not implement the changes that were necessary. As far as resource allocation was concerned, Nokia saw itself primarily as a hardware company rather than a software company. Its engineers were great at designing and producing hardware, but not the programs that drive the devices. They underestimated the importance of software (including, crucially, the apps that run on smartphones). Largely it was hardware rather than software experts who controlled its development process. By contrast, Apple had always emphasised that hardware and software were equally important. Yet while they were losing their dominance, Nokia were well aware of most of the changes occurring in the mobile communications market and the technology developments being actively pursued by competitors. Arguably, they were not short of awareness, but they did lack the capacity to convert awareness into action. The failure of big companies to adapt to changing circumstances is one of the fundamental puzzles in the world of business, says Professor Birkinshaw. Occasionally, a genuinely 'disruptive' technology can wipe out an entire industry. But usually the sources of failure are less dramatic. Often it is a failure to implement strategies or technologies that have already been developed, an arrogant disregard for changing customer demands, or a complacent attitude towards new competitors.

DIAGNOSTIC QUESTION

Does operations strategy set an improvement path?

An operations strategy is the starting point for operations improvement. It sets the direction in which the operation will change over time. It is implicit that the business will want operations to change for the better. Therefore, unless an operations strategy gives some idea as to how improvement will happen, it is not fulfilling its main purpose. This is best thought about in terms of how performance objectives, both in themselves and relative to each other, will change over time. To do this, we need to understand the concept of, and the arguments concerning, the trade-offs between performance objectives.

An operations strategy should guide the trade-offs between performance objectives

An operations strategy should address the relative priority of operation's performance objectives ('for us, speed of response is more important than cost efficiency, quality is more important than variety', and so on). To do this it must consider the possibility of improving its performance in one objective by sacrificing performance in another. So, for example, an operation might wish to improve its cost efficiencies by reducing the variety of products or services that it offers to its customers. Taken to its extreme; this 'trade-off' principle implies that improvement in one performance objective can *only* be gained at the expense of another. 'There is no such thing as a free lunch' could be taken as a summary of this approach to managing. Probably the best-known summary of the trade-off idea comes from Professor Wickham Skinner, the most influential of the originators of the strategic approach to operations, who said:

> *. . . most managers will readily admit that there are compromises or trade-offs to be made in designing an airplane or truck. In the case of an airplane, trade-offs would involve matters such as cruising speed, take-off and landing distances, initial cost, maintenance, fuel consumption, passenger comfort and cargo or passenger capacity. For instance, no one today can design a 500-passenger plane that can land on an aircraft carrier and also break the sound barrier. Much the same thing is true in . . . [operations].*[17]

OPERATIONS PRINCIPLE

In the short term, operations cannot achieve outstanding performance in all its operations objectives simultaneously.

But there is another view of the trade-offs between performance objectives. This sees the very idea of trade-offs as the enemy of operations improvement, and regards the acceptance that one type of performance can only be achieved at the expense of another as both limiting and unambitious. For any real improvement of total performance, it holds, the effect of trade-offs must be overcome in some way. In fact, overcoming trade-offs must be seen as the central objective of strategic operations improvement.

These two approaches to managing trade-offs result in two approaches to operations improvement. The first approach emphasises 'repositioning' performance objectives by trading-off improvements in some objectives for a reduction in performance in other; the second one emphasises increasing the 'effectiveness' of the operation by overcoming trade-offs so that improvements in one or more aspects of performance can be achieved without any reduction in the performance of others. Most businesses at some time or other will adopt both approaches. This is best illustrated through the concept of the 'efficient frontier' of operations performance.

OPERATIONS PRINCIPLE

In the long term, a key objective of operations strategy is to improve all aspects of operations performance.

Trade-offs and the efficient frontier

Figure 2.12(a) shows the relative performance of several companies in the same industry in terms of their cost efficiency and the variety of products or services that they offer to their customers. Presumably all the operations would ideally like to be able to offer very high variety, while still having very high levels of cost efficiency. However, the increased complexity that a high variety of product or service offerings brings will generally reduce the operation's ability to operate efficiently. Conversely, one way of improving cost efficiency is to severely limit the variety on offer to customers. The spread of results in Figure 2.12(a) is typical of an exercise such as this. Operations A, B, C and D all have chosen a different balance between variety and cost efficiency. But none is dominated by any other operation in the sense that another operation necessarily has 'superior' performance. Operation X, however, has an inferior performance because operation A

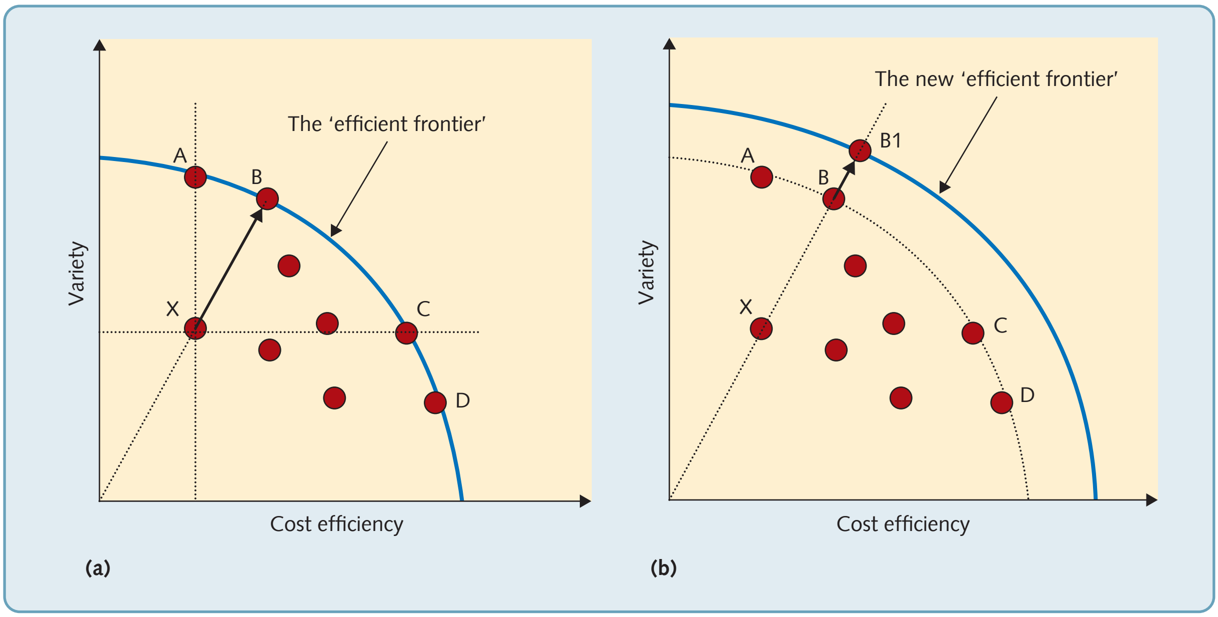

Figure 2.12 The efficient frontier

is able to offer higher variety at the same level of cost efficiency; and operation C offers the same variety but with better cost efficiency. The convex line on which operations A, B, C and D lie is known as the 'efficient frontier'. They may choose to position themselves differently (presumably because of different market strategies) but they cannot be criticised for being ineffective. Of course any of these operations that lie on the efficient frontier may come to believe that the balance they have chosen between variety and cost efficiency is inappropriate. In these circumstances they may choose to reposition themselves at some other point along the efficient frontier.

OPERATIONS PRINCIPLE

Operations that lie on the 'efficient frontier' have performance levels that dominate those which do not.

By contrast, operation X has also chosen to balance variety and cost efficiency in a particular way but is not doing so effectively. Operation B has the same ratio between the two performance objectives but is achieving them more effectively. Operation X will generally have a strategy that emphasises increasing its effectiveness before considering any repositioning.

However, a strategy that emphasises increasing effectiveness is not confined to those operations that are dominated, such as operation X. Those with a position on the efficient frontier will generally also want to improve their operation's effectiveness by overcoming the trade-off that is implicit in the efficient frontier curve. For example, suppose operation B in Figure 2.12(b) is the metrology company described earlier in this chapter. By adopting a modular product design strategy it improved both its variety and its cost efficiency simultaneously (and moved to position B1). What has happened is that operation B has adopted a particular operations practice (modular design) that has pushed out the efficient frontier. This distinction between positioning on the efficient frontier and increasing operations effectiveness to reach the frontier is an important one. Any operations strategy must make clear the extent to which it is expecting the operation to reposition itself in terms of its performance objectives and the extent to which it is expecting the operation to improve its effectiveness.

Improving operations effectiveness by using trade-offs

Improving the effectiveness of an operation by pushing out the efficient frontier requires different approaches, depending on the original position of the operation on the frontier. For example, in Figure 2.13 operation P is originally resourced and designed to offer a high level of

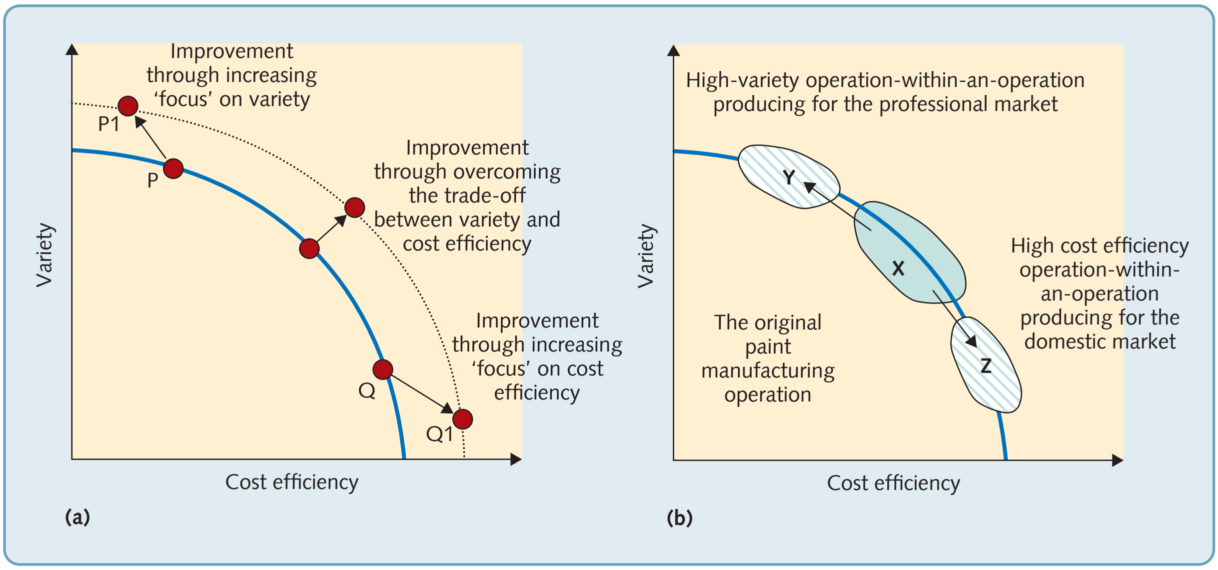

Figure 2.13 **Operations 'focus' and the 'operation-within-an-operation' concept illustrated using the efficient frontier model**

variety of products or services at the expense of low-cost efficiency. It has probably reached this position by adopting a series of operations practices that enable it to offer the variety, even if these practices are intrinsically expensive. For example, it may have invested in general-purpose technology and recruited employees with a wide range of skills. Improving variety even further may mean adopting more extreme operations practices that emphasise variety. For instance, it may reorganise its processes so that each of its larger customers has a dedicated set of resources that understands the specific requirements of that customer and can organise itself to totally customise every product or service it produces. This will probably mean a further sacrifice of cost efficiency, but it allows an ever greater variety of products or services to be produced (P1). Similarly, operation Q may increase the effectiveness of its cost efficiency, by becoming even less able to offer any kind of variety (Q1). For both operations P and Q, effectiveness is being improved through increasing the focus of the operation on a very narrow set of performance objectives and accepting an even further reduction in other aspects of performance.

The same principle of focus also applies to organisational units smaller than a whole operation. For example, individual processes may choose to position themselves on a highly focused set of performance objectives that match the market requirements of their own customers. So, for example, a business that manufactures paint for interior decoration may serve two quite distinct markets. Some of its products are intended for domestic customers who are price sensitive but demand only a limited variety of colours and sizes. The other market is professional interior decorators that demand a very wide variety of colours and sizes but are less price sensitive. The business may choose to move from a position where all types of paint are made on the same processes (position X in Figure 2.13(b)) to one where it has two separate sets of processes (Y and Z); one that only makes paint for the domestic market and the other that only makes paint for the professional market. In effect, the business has segmented its operations processes to match the segmentation of the market. This is called the 'operation-within-an-operation' (or 'plant-within-a-plant', or 'shop-within-a-shop', etc.) concept.

Improving operations effectiveness by overcoming trade-offs

This concept of highly focused operations is not universally seen as appropriate. Many companies attempt to give 'the best of both worlds' to their customers. At one time, for example, a high-quality, reliable and error-free automobile was inevitably an expensive automobile. Now, with few exceptions, we expect even budget-priced automobiles to be reliable and almost free of any defects. Auto manufacturers found that not only could they reduce the number of defects on their vehicles without necessarily incurring extra costs, but also they could actually reduce costs by reducing errors in manufacture. If auto manufacturers had adopted a purely focused-based approach to improvement over the years, we may now only be able to purchase either very cheap low-quality automobiles or very expensive high-quality automobiles. So a permanent expansion of the efficient frontier is best achieved by overcoming trade-offs through improvements in operations practice.

> **OPERATIONS PRINCIPLE**
> *An operation's strategy improvement path can be described in terms of repositioning and/or overcoming its performance trade-offs.*

Even trade-offs that seem to be inevitable can be reduced to some extent. For example, one of the decisions that any supermarket manager has to make is how many checkout positions to open at any time. If too many checkouts are opened then there will be times when the checkout staff do not have any customers to serve and will be idle. The customers, however, will have excellent service in terms of little or no waiting time. Conversely, if too few checkouts are opened, the staff will be working all the time but customers will have to wait in long queues. There seems to be a direct trade-off between staff utilisation (and therefore cost) and customer waiting time (speed of service). Yet even the supermarket manager deciding how many checkouts to open can go some way to affecting the trade-off between customer waiting time and staff utilisation. The manager might, for example, allocate a number of 'core' staff to operate the checkouts but also arrange for those other staff who are performing other jobs in the supermarket to be trained and 'on-call' should demand suddenly increase. If the manager on duty sees a build-up of customers at the checkouts, these other staff could quickly be used to staff checkouts. By devising a flexible system of staff allocation, the manager can both improve customer service and keep staff utilisation high.

Critical commentary

- Starting any discussion of strategy from a stakeholder perspective is far from undisputed. Read this'. 'At the economy-wide or social level, the issue is this: If we could dictate the criterion or objective function to be maximised by firms (and thus the performance criterion by which corporate executives choose among alternative policy options), what would it be? Or, to put the issue even more simply: How do we want the firms in our economy to measure their own performance? How do we want them to determine what is better versus worse?' who also holds that using stakeholder perspectives gives undue weight to narrow special interests who want to use the organisation's resources for their own ends. The stakeholder perspective gives them a spurious legitimacy which 'undermines the foundations of value-seeking behaviour'.[18]

- Similarly, the idea that operations strategy could ever become the driver of a business's overall strategy, and the associated concept of the resource-based view of the firm, are both problematic to some theorists. Business strategies, and functional strategies, were, for many years, seen as, first, market driven and second, planned in a systematic and deliberative manner. So, it became almost axiomatic to see strategy as starting from a full understanding of market positioning. In fact, the main source of sustainable competitive advantage was seen as unequivocally associated with how a business positioned itself in its markets. Get the market proposition right and customers would respond by giving you business. Get it wrong and they would go to the competitors with a better offering. Strategy was seen as aligning the whole organisation to the market position that could achieve long-term profitable differentiation when compared to competitors. Functional strategies were simply a more detailed interpretation of this overall imperative. Furthermore, strategy must be something that could be planned and directed. If managers could not influence strategy, then how could business be anything other than a lottery?

- The idea that sustainable competitive advantage could come from the capabilities of one's resources was a clear threat to the established position. Furthermore, the idea that strategies emerged, sometimes haphazardly and unpredictably, over time rather than were deliberate decisions taken by senior managers was also seemingly counterintuitive. Yet there is now considerable research evidence to support both these, once outrageous, propositions. Yet widely practised approaches to developing operations strategies are still ignoring (or downplaying) the idea of resourced-based strategic advantage. For example, the business-model/operating-model idea is predominantly a 'top-down' philosophy that relegates operations capabilities to a 'supporting' rather than 'driving' role. The position we have taken in this chapter is one of blending some aspects of the traditional view with the more recent ideas. Nevertheless, it is important to understand that there are still different views on the very nature of strategic management.

SUMMARY CHECKLIST

Each chapter contains a summary checklist in the form of questions that can be usefully applied to any type of operations and reflect the major diagnostic questions used within the chapter.

- ☐ Does the operation have a fully articulated operations strategy?
- ☐ Does it include a vision for the role and contribution of the operations function?
- ☐ Does operations strategy take significant stakeholders into account?
- ☐ What position on the Hayes and Wheelwright Stage 1 to 4 Model are your operations?
- ☐ Is there a recognised process for translating business strategy 'top-down' into operations strategy?
- ☐ Does operations strategy demonstrate both correspondence and coherence with business strategy?
- ☐ Does the organisation's business model fit with its operating model?
- ☐ Are the operation's performance objectives fully articulated?
- ☐ Are performance objectives understood in terms of whether they are order-winners or qualifiers?
- ☐ Have different parts of the operation (probably producing different products or services) their own relative priority of performance objectives that reflect their possibly different competitive positions?
- ☐ Is there a recognised process for bottom-up communication on strategic issues?
- ☐ Are the main strategic decisions that shape operation's' resources fully identified?
- ☐ Is the idea of operations-based capabilities fully understood?
- ☐ What capabilities does the operation currently possess?
- ☐ Are these operations and/or resources, scarce, imperfectly mobile, imperfectly imitable, or imperfectly substitutable?
- ☐ Are the logical links established between what the market requires (in terms of performance objectives) and what capabilities an operations possesses (in terms of the major strategic decision areas)?
- ☐ Where would you put the operation in terms of Figure 2.11 that describes the line of fit between market requirements and operations capabilities?
- ☐ Have the key trade-offs for the operation been identified?
- ☐ What combination of repositioning, in order to change the nature of trade-offs, and overcoming the trade-offs themselves, is going to be used to improve overall operations performance?

CASE STUDY

McDonald's: half a century of growth[19]

It's loved and it's hated. It is a shining example of how good value food can be brought to a mass market. It is a symbol of everything that is wrong with 'industrialised', capitalist, bland, high-calorie, and environmentally unfriendly commercialism. It is the best-known and most loved fast food brand in the world, with more than 32,000 restaurants in 117 countries, providing jobs for 1.7 million staff and feeding 60 million customers per day (yes, per day!). It is part of the homogenisation of individual national cultures, filling the world with bland, identical; 'cookie cutter', Americanised and soulless operations that dehumanise its staff by forcing them to follow ridged and over-defined procedures. But whether you see it as friend, foe, or a bit of both, McDonald's has revolutionised the food industry, affecting the lives both of the people who produce food and the people who eat it. It has also had its ups (mainly) and downs (occasionally). Yet, even in the toughest times it has always displayed remarkable resilience. Even after the economic turbulence of 2008, McDonald's reported an exceptional year of growth in 2009 posting sales increases and higher market share around the world. It was the sixth consecutive year of positive sales in every geographic region of their business.

Starting small

Central to the development of McDonald's is Ray Kroc, who by 1954 and at the age of 52 had been variously a piano player, a paper cup salesman and a multi-mixer salesman. He was surprised by a big order for eight multi-mixers from a restaurant in San Bernardino, California. When he visited the customer he found a small but successful restaurant run by two brothers Dick and Mac McDonald. They had opened their 'Bar-B-Que' restaurant 14 years earlier adopting the usual format at that time; customers would drive-in, choose from a large menu and be served by a 'car hop'. However, by the time Ray Kroc visited the Brothers' operation it had changed to a self-service drive-in format with a limited menu of nine-items. He was amazed by the effectiveness of their operation. Focusing on a limited menu including burgers, fries and beverages, had allowed them to analyse every step of the process of producing and serving their food. Ray Kroc was so overwhelmed by what he saw that he persuaded the brothers to adopt his vision of creating McDonald's restaurants all over the U.S, the first of which opened in Des Plaines, Illinois in June 1955. However, later, Kroc and the McDonald brothers 'quarrelled, and Kroc bought the brothers out. Now with exclusive rights to the McDonald's name, the restaurants spread, and in 5 years there were 200 restaurants through the US. After 10 years the company went public; the share price doubling in the first month. But through this, and later, expansion, Kroc insisted on maintaining the same principles that he had seen in the original operation. '*If I had a brick for every time I've repeated the phrase Quality, Service, Cleanliness and Value, I think I'd probably be able to bridge the Atlantic Ocean with them.*' (Ray Kroc)

Priority to the process

Ray Kroc had been attracted by the cleanliness, simplicity, efficiency and profitability of the McDonald brothers' operation. They had stripped fast-food delivery down to its essence and eliminated needless effort to make a swift assembly line for a meal at a reasonable price. Kroc wanted to build a process that would become famous for food of consistently high quality, using uniform methods of preparation. His burgers, buns, fries and beverages should taste just the same in Alaska as they did in Alabama. The answer was the 'Speedee Service System'; a standardised process that prescribed exact preparation methods, specially designed equipment and strict product specifications. The emphasis on process standardisation meant that customers could be assured of identical levels of food and service quality every time they visited any store, anywhere. Operating procedures were specified in minute detail. In its first operations manual, which by 1991 had reached 750 pages, it prescribed specific cooking instructions such as temperatures, cooking times and portions to be followed rigorously. Similarly, operating procedures were defined to ensure the required customer experience; for example, no food items were to be held more than 10 minutes in the transfer bin between being cooked and being served. Technology was also automated. Specially designed equipment helped to guarantee consistency using 'fool-proof' devices. For example; the ketchup was dispensed through a metered pump. Specially designed 'clam shell' grills cooked both sides of each meat patty simultaneously for a pre-set time. And when it became clear that the metal tongs used by staff to fill French-fry containers were awkward to use efficiently, McDonald's engineers devised a simple V-shaped aluminium scoop that made the job faster and easier as well as presenting the fries in a more attractive alignment with their container.

For Kroc, the operating process was both his passion and the company's central philosophy. It was also the foundation of learning and improvement. The company's almost compulsive focus on process detail was not an end in itself. Rather it was to learn what contributed to consistent high-quality service in practice and what did not. Learning was always seen as important by McDonalds. In 1961, it founded 'Hamburger University', initially in the basement of a restaurant in Elk Grove Village, Illinois. It had a research and development laboratory to develop new cooking, freezing, storing and serving methods. Also franchisees and operators were trained in the analytical techniques necessary to run a successful McDonald's. It awarded degrees in 'Hamburgerology'. But learning was not just for headquarters. The company also formed a 'field service' unit to appraise and help its restaurants by sending field service consultants to review their performance on a number of 'dimensions' including cleanliness, queuing, food quality and customer service. As Ray Kroc, said, '*We take the hamburger business more seriously than anyone else. What sets McDonald's apart is the passion that we and our suppliers share around producing and delivering the highest-quality beef patties. Rigorous food safety and quality standards and practices are in place and executed at the highest levels every day*.'

No story illustrates the company's philosophy of learning and improvement better than its adoption of frozen fries. French fried potatoes had always been important. Initially, the company tried observing the temperature levels and cooking methods that produced the best fries. The problem was that the temperature during the cooking process was very much influenced by the temperature of the potatoes when they were placed into the cooking vat. So, unless the temperature of the potatoes before they were cooked were also controlled (not very practical) it was difficult to specify the exact time and temperature that would produce perfect fries. But McDonald's researchers have perseverance. They discovered that, irrespective of the temperature of the raw potatoes, fries were always at their best when the oil temperature in the cooking vat increased by three degrees above the low temperature point after they were put in the vat. So by monitoring the temperature of the vat, perfect fries could be produced every time. But that was not the end of the story. The ideal potato for fries was the Idaho Russet, which was seasonal and not available in the summer months, when an alternative (inferior) potato was used. One grower, who, at the time, supplied a fifth of McDonald's potatoes, suggested that he could put Idaho Russets into cold storage for supplying during the summer period. Notwithstanding investment in cold storage facilities, all the stored potatoes rotted. Not to be beaten, he offered another suggestion. Why don't McDonald's consider switching to frozen potatoes? This was no trivial decision and the company was initially cautious about meddling with such an important menu item. However, there were other advantages in using frozen potatoes. Supplying fresh potatoes in perfect condition to McDonald's rapidly expanding chain was increasingly difficult. Frozen potatoes could actually increase the quality of the company's fries if a method of satisfactorily cooking them could be found. Once again McDonald's developers came to the rescue. They developed a method of air drying the raw fries, quick frying, and then freezing them. The supplier, who was a relatively small and local supplier when he first suggested storing Idaho Russets, grew its business to supply around half of McDonalds US business.

Throughout their rapid expansion a significant danger facing McDonald's was losing control of their operating system. They avoided this, partly by always focusing on four areas: improving the product; establishing strong supplier relationships; creating (largely customised) equipment; and developing franchise holders. But also it was their strict control of the menu which provided a platform of stability. Although their competitors offered a relatively wide variety of menu items, McDonalds limited theirs to ten items. This allowed uniform standards to be established, which in turn encouraged specialisation. As one of McDonald's senior managers at the time stressed, '*It wasn't because we were smarter. The fact that we were selling just ten items [and,] had a facility that was small, and used a limited number of suppliers created an ideal environment*.' Capacity growth (through additional stores) was also managed carefully. Well-utilised stores were important to franchise holders, so franchise opportunities were located only where they would not seriously undercut existing stores. Ray Kroc used the company plane to spot from the air the best locations and road junctions for new restaurant branches.

Securing supply

McDonald's said that it has been the strength of the alignment between the Company, its franchisees and its suppliers (collectively referred to as the System) that has been the explanation for its success. Expanding the McDonalds chain, especially in the early years meant persuading both franchisees and suppliers to buy into the company's vision: '*Working,' as Ray Kroc put it, 'not for McDonald's, but for themselves, together with McDonald's.' He promoted the slogan, 'In business for yourself, but not by yourself*.' But when they started, suppliers proved problematic. McDonald's approached the major food suppliers, such as Kraft and Heinz, but without much success. Large and established suppliers were reluctant to conform to McDonald's requirements, preferring to focus on retail sales. It was the relatively small companies who were willing to risk supplying what seemed then to be a risky venture. Yet, as McDonald's grew, so did its suppliers. Also, McDonald's relationship with its suppliers was seen as less adversarial than with some other customers. One supplier is quoted as saying; '*Other chains would walk away from you for half a cent. McDonald's was more concerned with getting quality. McDonald's always treated me with respect even when they became much bigger and didn't have to*.' Furthermore, suppliers were always seen as a source of innovation. For example, one of McDonald's meat suppliers, Keystone Foods, developed a novel quick-freezing process that captured the fresh taste and texture of beef patties. This meant that every patty could retain its consistent quality until it hit the grill. Keystone shared its technology with other McDonald's meat suppliers for McDonald's, and today the process is an industry standard. Yet, although innovative and close, supplier relationships were also rigorously controlled. Unlike some competitors who simply accepted what suppliers provided, complaining only when supplies were not up to standard, McDonald's routinely analysed its supplier's products.

Fostering franchisees

McDonald's revenues consist of sales by company-operated restaurants and fees from restaurants operated by franchisees. McDonalds view themselves primarily as a franchisor and believe franchising is . . . '*important to delivering great, locally-relevant customer experiences and driving profitability*'. However, they also believe that directly operating restaurants is essential to providing the company with real operations experience. In 2009, of the 32,478 restaurants in 117 countries, 26,216 were operated by franchisees and 6,262 were operated by the Company. Where McDonalds was different to other franchise operations was in their relationships. Some restaurant chains concentrated on recruiting franchisees that may then be ignored. McDonalds, on the other hand, expected its franchisees to contribute their experiences for the benefit of all. Ray Kroc's original concept was that franchisees would make money before the company did. So he made sure that the revenues that went to McDonalds came from the success of the restaurants themselves, rather than from initial franchise fees.

Initiating innovation

Ideas for new menu items have often come from franchisees. For example, Lou Groen, a Cincinnati franchise holder had noticed that in Lent (a 40-day period when some Christians give up eating red meat on Fridays and instead eat only fish or no meat at all) some customers avoided the traditional hamburger. He went to Ray Kroc, with his idea for a 'Filet-o-Fish®'; a steamed bun with a shot of tartar sauce, a fish fillet, and cheese on the bottom bun. But Kroc wanted to push his own meatless sandwich, called the hula burger; a cold bun with a piece of pineapple and cheese. Groen and Kroc competed on a Lenten Friday to see whose sandwich would sell more. Kroc's hula burger failed, selling only six sandwiches all day while Groen sold 350 Filet-o-Fish. Similarly, the Egg McMuffin® was introduced by franchisee Herb Peterson, who wanted to attract customers into his McDonalds stores all through the day, not just at lunch and dinner. He came up with idea for the signature McDonald's breakfast item because he was reputedly 'very partial to eggs Benedict and wanted to create something similar'.

Other innovations came from the company itself. By the beginning of the 1980s, poultry was becoming more fashionable to eat and sales of beef were sagging. Fred Turner, then the Chairman of McDonald's had an idea for a new meal; a chicken finger-food without bones, about the size of a thumb. After six months of research, the food technicians and scientists managed to reconstitute shreds of white chicken meat into small portions which could be breaded, fried, frozen then reheated. Test-marketing the new product was positive, and in 1983 they were launched under the name Chicken McNuggets®. These were so successful that within a month McDonald's became the second largest purchaser of chicken in the USA. By 1992, Americans were eating more chicken than beef.

Other innovations came as a reaction to market conditions. Criticised by nutritionists, who worried about calorie-rich burgers, and shareholders who were alarmed by flattening sales, McDonald's launched its biggest menu revolution in 30 years in 2003 when it entered the prepared salad market. They offered a choice of dressings for their grilled chicken salad with Caesar dressing (and croutons) or the lighter option of a drizzle of balsamic dressing. Likewise, recent moves towards coffee sales were prompted by the ever-growing trend set by big coffee shops like Starbucks. McCafé®, a coffee-house-style food and drink chain, owned by McDonald's, had expanded to about 1,300 stores worldwide by 2011.

Problematic periods

The period from the early 1990s to the mid-2000s was difficult for parts of the McDonalds empire. Although growth in many parts of the world continued, in some developed markets, the company's hitherto rapid growth stalled. Partly this was due to changes in food fashion, nutritional concerns and demographic changes. Partly it was because competitors were learning to either emulate McDonald's operating system, or focus on one aspect of the traditional 'quick service' offering, such as speed of service, range of menu items, (perceived) quality of food, or price. Burger King, promoted itself on its 'flame-grilled' quality. Wendy's offered a fuller service level. Taco Bell undercut McDonalds' prices with their 'value pricing' promotions. Drive-through specialists such as Sonic speeded up service times. But it was not only competitors that were a threat to McDonald's growth. So-called 'fast food' was developing a poor reputation in some quarters and, as its iconic brand, McDonald's was taking much of the heat. Similarly, the company became a lightning rod for other questionable aspects of modern life that it was held to promote, from cultural imperialism, low-skilled jobs, abuse of animals, the use of hormone-enhanced beef, to an attack on traditional (French) values (in France). A French farmer called Jose Bové (who was briefly imprisoned) got other farmers to drive their tractors through, and wreck, a half-built McDonald's. When he was tried, 40,000 people rallied outside the courthouse.

The Chief Executive of McDonald's in the UK, Jill McDonald (yes, really!), said that some past difficulties were self-induced. They included a refusal to face criticisms and a reluctance to acknowledge the need for change. '*I think by the end of 1990s we were just not as close to the customer as we needed to be, we were given a hard time in the press and we lost our confidence. We needed to reconnect, and make changes that would disrupt people's view of McDonald's.' Investing in its people also needed to be re-emphasised. 'We invest about £35m a year in training people. We have become much more of an educator than an employer of people'. Nor does she accept the idea of 'McJobs' (meaning boring, poorly paid, often temporary jobs with few prospects). 'That whole McJob thing makes me so angry. It's snobbish. We are the biggest employer of young people in Britain. Many join us without qualifications. They want a better life, and getting qualifications is something they genuinely value.*'

Surviving strategies

Yet, in spite of its difficult period, the company has not only survived, but through the late-2000s has thrived. In 2009 McDonald's results showed that in the US, sales and market share both grew for the seventh consecutive year, with new products such as McCafé premium coffees, the premium Angus Third Pounder, smoothies and frappes, together with more convenient locations, extended hours, efficient drive-thru service and value-oriented promotions. In the UK, changes to the stores' decor and adapting menus have also helped stimulate growth. Jill McDonald's views are not untypical of other regions, '*We have probably changed more in the past four years than the past 30: more chicken, 100% breast meat, snack wraps, more coffee – lattes and cappuccinos, ethically sourced, not at rip-off prices. That really connected with customers. We sold 100m cups last year*.'

Senior managers put their recent growth down to the decision in 2003 to reinvent McDonald's by becoming '*better, not just bigger' and implementing its 'Plan to Win' This focused on 'restaurant execution', with the goal of . . . 'improving the overall experience for our customers'. It provided a common framework for their global business, yet allowed for local adaptation. Multiple improvement initiatives were based on its 'five key drivers of exceptional customer experiences*' (People, Products, Place, Price and Promotion). But what of McDonald's famous standardisation? During its early growth no franchise holder could deviate from the 700+ page McDonald's operations manual known as 'the Bible'. Now things are different, at least partly because different regions have developed their own products. In India, the 'Maharaja Mac' is made of mutton, and the vegetarian options contain no meat or eggs. Similarly, McDonald's in Pakistan offers three spicy 'McMaza meals'. Even in the USA things have changed. In at least one location in Indiana, there's now a McDonald's with a full service 'Diner' inside, where waitresses serve 100 combinations of food, on china; a far cry from Ray Kroc's vision of stripping out choice to save time and money.'

QUESTIONS

1. How has competition to McDonald's changed over time?
2. What are the main operations performance objectives for McDonald's?
3. Draw an operations strategy matrix for McDonald's.

APPLYING THE PRINCIPLES

Some of these exercises can be answered by reading the chapter. Others will require some general knowledge of business activity and some might require an element of investigation. **Model answers for the first two questions are available on the student companion website.** For model answers to the other questions in this section, please ask your tutor/lecturer.

1 The environmental services department of a city has two recycling services – newspaper collection (NC) and general recycling (GR). The NC service is a door-to-door collection service that, at a fixed time every week, collects old newspapers that householders have placed in reusable plastic bags at their gate. An empty bag is left for the householders to use for the next collection. The value of the newspapers collected is relatively small; the service is offered mainly for reasons of environmental responsibility. By contrast the GR service is more commercial. Companies and private individuals can request a collection of materials to be disposed of, either using the telephone or the internet. The GR service guarantees to collect the material within 24 hours unless the customer prefers to specify a more convenient time. Any kind of material can be collected and a charge is made depending on the volume of material. This service makes a small profit because the revenue both from customer charges and from some of the more valuable recycled materials exceeds the operation's running costs. How would you describe the differences between the performance objectives of the two services?

2 The Managing Partner of The Branding Partnership (TBP) describes her business, *'It is about four years now since we specialised in the small-to-medium firms' market. Before that we also used to provide brand consultancy services for anyone who walked in the door. So now we have built up our brand consultancy skills in many areas. However, within the firm, I think we could focus our activities even more. There seem to be two types of assignment that we are given. About 40 per cent of our work is relatively routine. Typically, these assignments are conventional market research and focus group exercises. Both these activities involve a standard set of steps that can be carried out by comparatively junior staff. Of course, an experienced consultant is needed to make some decisions; however, most of this work is fairly straightforward. Customers expect us to be relatively inexpensive and fast in delivering the service. They do not expect us to make simple errors; in fact if we did this too often we would lose business. Fortunately, our customers know that they are buying a 'standard package' and don't expect it to be too customised. The problem here is that specialist agencies have been emerging over the last few years and they are starting to undercut us on price. Yet, I still feel that we can operate profitably in this market and anyway, we still need these capabilities to serve our other clients. The other 60 per cent of our work is for clients who require far more specialist services, such as assignments involving major brand reshaping. These assignments are complex, large, take longer, and require significant branding skill and judgment. It is vital that clients respect and trust the advice we give them in all 'brand associated' areas such as product development, promotion, pricing, and so on. Of course they assume that we will not be slow or unreliable in preparing advice, but mainly it's trust in our judgment backed up by hard statistics that is important to the client. This is popular work with our staff. It is both interesting and very profitable.'* How different are the two types of business described by the Managing Partner of TBP? It has been proposed that she split the firm into two separate businesses; one to deal with routine services and the other to deal with more complex services. What would be the advantages and disadvantages of doing this?

3 DSD designs, makes and supplies medical equipment to hospitals and clinics. Its success was based on their research and development culture. Although around 50 per cent of manufacturing was done in-house, their products were relatively highly priced, but customers were willing to pay for their technical excellence and willingness to customise equipment. Around 70 per cent of all orders involved some form of customisation from standard 'base models'. Manufacturing could take 3 months from receiving the specification to completing assembly, but customers were more interested in equipment being delivered on time rather than immediate availability. According to their CEO, *'manufacturing is really a large laboratory. The laboratory-like culture helps us to maintain our superiority in leading edge product technology and customisation. It also means*

that we can call upon our technicians to pull out all the stops in order to maintain delivery promises. However, I'm not sure how manufacturing, or indeed the rest of the company, will deal with the new markets and products which we are getting into'.

The new products were 'small black box' products that the company had developed. These were devices that could be attached to patients, or implanted. They took advantage of sophisticated electronics and could be promoted directly to consumers as well as to hospitals and clinics. The CEO knew their significance. *'Although expensive, we have to persuade health care and insurance companies to encourage these new devices. More problematic is our ability to cope with these new products and new markets. We are moving towards being a consumer company, making and delivering a higher volume of more standardised products where the underlying technology is changing fast. We must become faster in our product development. Also, for the first time, we need some kind of logistics capability. I'm not sure whether we should deliver products ourselves or subcontract this. Manufacturing faces a similar dilemma. On one hand it is important to maintain control over production to ensure high quality and reliability; on the other hand, investing in the process technology to make the products will be very expensive. There are subcontractors who could manufacture the products; they have experience in this kind of manufacturing but not in maintaining the levels of quality we will require. We will also have to develop a 'demand fulfillment' capability to deliver products at short notice. It is unlikely that customers would be willing to wait the 3 months our current customers tolerate. Nor are we sure of how demand might grow. I'm confident that growth will be fast but we will have to have sufficient capacity in place not to disappoint our new customers. We must develop a clear understanding of the new capabilities that we will have to develop if we are to take advantage of this wonderful market opportunity.'*

What advice would you give DSD? Consider the operational implication of entering this new market.

4 Xexon7 is a specialist artificial intelligence (AI) development firm that develops algorithms for various on-line services. As part of its client services it has a small (10 person) help-desk call centre to answer client queries. Clients could contact them from anywhere in the world at any time of the day or night with a query. Demand at any point in time is fairly predictable, especially during the (European) daytime. Demand during the night hours (Asia and the Americas) is considerably lower than in the daytime and also less predictable. *'Most of the time we forecast demand pretty accurately, so we can schedule the correct number of employees to staff the work stations. There is still some risk, of course. Scheduling too many staff at any point in time will waste money and increase our costs, while scheduling too few will reduce the quality and response of the service we give.'* (Peter Fisher, Help Desk Manager) Peter was, overall, pleased with the way in which his operation worked. However, he felt that a more systematic approach could be taken to identifying improvement opportunities. *'I need to develop a logical approach to identify how we can invest in improving things like sophisticated diagnostic systems. We need to both reduce our operating costs and maintain, and even improve, our customer service.'*

What are the trade-offs that must be managed in this type of call centre?

Notes on chapter

1 For a more thorough explanation, see Slack, N. and Lewis, M. (2017) *Operations Strategy* (5th edn), Pearson.

2 Hayes, R.H. and Wheelwright, S.C. (1984) *Restoring our Competitive Edge,* John Wiley.

3 For a more thorough explanation, *see* Slack, N. and Lewis, M. (2017) *Operations Strategy,* op.cit.

4 Based on an idea in Slack, N. (2017) *The Operations Advantage,* Kogan Page.

5 Osterwalder, A., Pigneur, Y. and Tucci, C. (2005) 'Clarifying business models: origins, present and future of the concept', *CAIS,* vol. 15, 751–775.'

6 Osterwalder, A.(2005) 'What is a business model?', http://business-model-design.blogspot.com/2005/11/what-is-business-model.html

7 Based on the definitions developed by Cap Gemini.

8 All quotes taken from each company's website.

9 This terminology was first proposed by Professor Terry Hill.

10 Mintzberg, H. and Waters, J.A. (1995) 'Of strategies: deliberate and emergent', *Strategic Management Journal,* July/Sept.

11 For a full explanation of this concept, see Slack, N. and Lewis, M. (2017) *Operations Strategy,* op.cit.

12 For an alternative categorisation of strategic decision areas, see Slack, N. and Lewis, M.A. (2017) *Operations Strategy,* op.cit.

13 An idea proposed by Jay Barney. See Barney, J.B. (2001) 'Is the resource-based theory a useful perspective for strategic management research? Yes' , *Academy of Management Review*.

14 Barney, J. B., & Hesterly, W. S. (2010), 'VRIO Framework', In *Strategic Management and Competitive Advantage* (68–86) Pearson.

15 Barney, J. B. (1995) 'Looking inside for competitive advantage', *Academy of Management Executive,* 9, (4) no. 49–61.

16 Sources include: Birkinshaw, J. (2013) 'Why corporate giants fail to change', *CNN Money,* 8 May; Magazine, A. (2013) 'Two lessons learned from Nokia's downfall', Techwell.com, 24 October; Hessman, T. (2013) 'The road to failure: Nokia, Blackberry and . . . Apple, *Industry Week,* 6 September.

17 A point made initially by Skinner. Skinner, W. (1985) *Manufacturing: The Formidable Competitive Weapon,* John Wiley.

18 From a speech by Michael Jensen of Harvard Business School.

19 Sources include: Kroc, R. A. (1977) *Grinding it Out: The Making of McDonald's,* St. Martin's Press; Love, J. (1995) *McDonald's: Behind the Golden Arches,* Random House Publishing Group; www.aboutmcdonalds.com (2009); Davidson, A. (2011) 'So Mrs McDonald, would you like fries with that?', *Sunday Times,* 13 February; McDonalds Annual Report, 2009; Upton, D. (1992) *McDonalds Corporation Case Study,* Harvard Business School.

TAKING IT FURTHER

Barney, J.B. (2007) *Resource-Based Theory: Creating and Sustaining Competitive Advantage, Oxford University Press*. *Great on the 'Resource-Based View of the Firm'*

Hayes, R. (2006) 'Operations, Strategy, and Technology: Pursuing the Competitive Edge', *Strategic Direction, vol. 22, Issue 7, Emerald Group Publishing Limited*. *A summary of the subject from one (if not the) leading academics in the area.*

Hayes, R.H. and Pisano, G.P., (1994) 'Beyond World Class: The New Manufacturing Strategy', *Harvard Business Review, vol. 72, no.* 1. *Same as above.*

Hill, T. (2006) *Manufacturing Operations Strategy (3rd edn), Palgrave Macmillan*. *The descendant of the first non-US book to have a real impact in the area.*

Johnson, G., Whittington, R. and Scholes, K. (2011) *Exploring Strategy (9th edn), Pearson*. *There are many good books on strategy. This is one of the best.*

Slack, N. and Lewis, M. (2017) *Operations Strategy (5th edn), Pearson*. *What can we say – just brilliant!*

3

Product and service innovation

Introduction

Customers want innovation, and the businesses that provide it can succeed to a greater degree than their competitors that are content to offer the same things (or services) to their customers, created in the same old way. Companies such as Google, Amazon, Netflix, Nike, Airbnb, Apple and Dropbox are successful because they have refreshed the idea of what their markets want. Their products and services have been continually updated, altered and modified. Some changes are small, incremental adaptations to existing ways of doing things. Others are radical, major departures from anything that has gone before. The innovation activity is about successfully delivering change in its many different forms. Being good at innovation has always been important. What has changed in recent years is the sheer speed and scale of innovation in industries all over the world. Innovation processes are also increasingly complex, with inputs from different individuals and departments within an organisation, and increasingly from a wide variety of external sources. Increasingly operations mangers are expected to take a greater and more active part in service and product innovation. This chapter examines what is meant by innovation and why it matters; the general stages involved in bringing service and product offerings from concept to launch; and some of the resourcing and organisational considerations for innovation (see Figure 3.1).

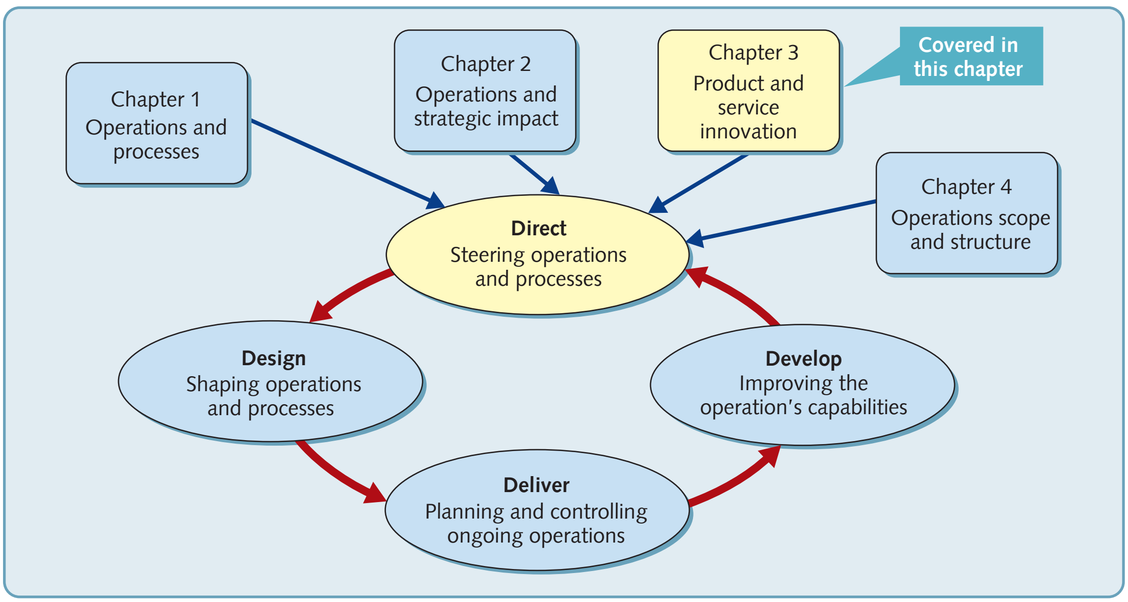

Figure 3.1 **Product and service innovation**

EXECUTIVE SUMMARY

What is the strategic role of product and service innovation?

↓

Are the product and service innovation process objectives specified?

↓

Is the product and service innovation process defined?

↓

Are the resources for product and service innovation adequate?

↓

Is the development of products and services and of the process that created them simultaneous?

What is the strategic role of product and service innovation?

A business's products and services are how markets judge it: they are its 'public face'. Innovation is all about doing things differently. That is why product and service innovation is so important. It involves changing what is offered to customers in some way, in addition to using 'creativity' and 'design', terms that have similar but different meanings. Creativity is the use of imagination or original ideas; it is an essential ingredient of innovation, which means doing something new. The design activity transforms innovation into a practical proposition by defining its looks or arrangement. Innovation often follows an S-shaped progress, where, in the early stages of the introduction of new ideas, relatively small performance improvements are experienced, but, with time, performance increases, and eventually levels off. A further strategic perspective on innovation is the difference between incremental and radical innovation. It is a distinction, together with that between 'architectural' and 'component' knowledge in the Henderson–Clark model. Yet, at its most basic, the innovation activity is a process that involves many of the same design issues common to other operations processes.

Are the product and service innovation process objectives specified?

The performance of the innovation process can be assessed in much the same way as we would consider the outputs from it. The quality of the innovation process can be judged in terms of both conformance (no errors in the offering) and specification (the effectiveness of the offering in achieving its market requirements). Speed in the innovation process is often called 'time to market' (TTM). Short TTM implies that the offering can be introduced to the market frequently, achieving strategic impact. Dependability in the innovation process means meeting launch delivery dates. In turn, this often requires that the innovation process is sufficiently flexible to cope with disruptions. The cost of the innovation process can be thought of both as the amount of budget that is necessary to develop a new offering and as the impact of the innovation on the cost of delivering the service or product on an ongoing basis. Finally, the sustainability of the innovation process considers the impact on broad stakeholder objectives encompassed within the 'triple bottom line' – people, planet and profit.

Is the product and service innovation process defined?

To create a fully specified service or product offering, potential designs tend to pass through a set of stages in the innovation process. Almost all stage models start with a general idea or 'concept' and progress through to a fully defined specification for the offering, incorporating various service and product components. In-between these two states, the offering may pass through the following stages: concept generation, concept screening, preliminary design (including consideration of standardisation, commonality, modularisation and mass customisation), evaluation and improvement, prototyping and final design.

Are the resources for developing product and service innovations adequate?

To be effective, the innovation process needs to be resourced adequately. The detailed principles of process design that are discussed in Chapters 5 and 6 are clearly applicable when developing new services or product offerings. However, because the innovation activity is often an operation in its own right, there are some more strategic issues to consider: how much capacity to devote to innovation, how much of the innovation activity to outsource and what kinds of technology to use in the development of new offerings.

Is the development of products and services and of the process that created them simultaneous?

The outputs from the innovation, in the form of new service and product offerings, are important inputs into the processes that create and deliver them on an ongoing basis. Therefore, it is often best to consider these in parallel rather than (as more traditionally done) in sequence. Merging the innovation process for new service and product offerings with the processes that create them is sometimes called simultaneous (or interactive) design. Its key benefit is the reduction in the time taken for the whole innovation activity. In particular, four simultaneous design factors can be identified that promote fast time to market. These are: routinely integrating the design of the product-service offering and the design of the process used to create and deliver them; overlapping the stages in the innovation process; the early deployment of strategic decision-making to resolve design conflict; and an organisational structure that reflects the nature of the offering.

DIAGNOSTIC QUESTION

What is the strategic role of product and service innovation?

A business's products and services are its 'public face'. They are what markets judge a company on: good products and services equals a good company. This is why it is important to devote time and effort to the way in which innovative ideas are incorporated in new products and services. Moreover, it has long been accepted that there is a connection between how companies go about developing innovative products and services and how successful those products and services are in the marketplace.

Innovation, design and creativity

Developing new products and services is a creative, and often innovative, process. What is the relationship between terms such as 'innovation', 'creativity' and 'design'? They have similar but different meanings, overlap to some extent and are clearly related to each other. It is best to start with what exactly it is that we mean by 'innovation'. In fact, there are many definitions. The term is notoriously ambiguous and lacks either a single definition or measure. It is *'...a new method, idea, product, etc.'* (*Oxford English Dictionary*), *'change that creates a new dimension of performance'* (Peter Drucker, a well-known management writer), *'the act of introducing something new'* (the *American Heritage Dictionary*), *'a new idea, method or device'* (*Webster Online Dictionary*), *'new knowledge incorporated in products, processes and services'*.[1] What runs through all these definitions is the idea of novelty and change. Innovation is simply about doing something new. However, it is worth noting that the idea of innovation is both broader and more complete than that of 'invention'. An 'invention' is also something that is novel or unique (usually applied to a device or method), but it does not necessarily imply that the novel device or method has the potential to be practical, economic, or capable of being developed commercially. Innovation goes further than 'invention'. It implies not just the novel idea, but also the process of transforming the idea into something that provides a return for an organisation's customers, owners, or both. The study of innovation, including what influences it and how to manage it, is a huge subject. However, creativity is one particularly important attribute that is central to innovation. 'Creativity' is the ability to move beyond conventional ideas, rules, or assumptions, in order to generate significant new ideas. It is a vital ingredient in innovation.

So, if creativity is an essential ingredient of innovation, and innovation implies making novel ideas into practical, commercial form, what is the process that transforms innovative ideas into something more concrete? It is 'design'. Innovation creates the novel idea; design makes it work in practice. Design, is to 'conceive the looks, arrangement, and workings of something'. A design must deliver a solution that will work in practice. Design is also an activity that can be approached at different levels of detail. One may envisage the general shape and intention of something before getting down to defining its details (we shall observe this later in this chapter when we examine the process of product and service design). Figure 3.2 illustrates the relationship between creativity, innovation, and design as we use the terms here. These concepts are intimately related, which is why we treat them in the same chapter. First, we will look at some of the basic ideas that help to understand innovation.

Innovation is all about doing things differently. This involves working with customers, suppliers and internal functions to generate new ideas and then successfully implement them. Contributions are needed from those who understand market requirements; those

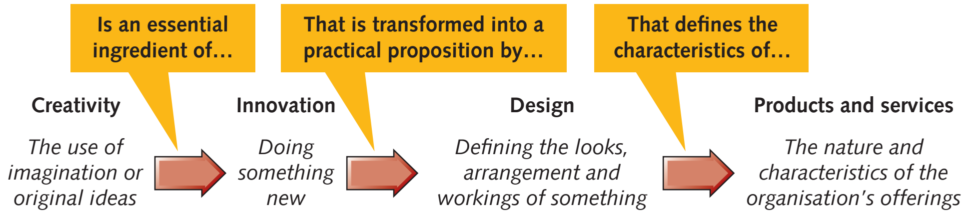

Figure 3.2 The relationship between creativity, innovation and design

who understand the technical aspects of the offering; those with access to cost and investment information; those who can protect intellectual property; and, most importantly, the operations people responsible for the innovation's delivery. Nevertheless, innovation depends on more than simply having good ideas. Look at the two contrasting examples that follow.

EXAMPLE

How the iPhone disrupted the smartphone market[2]

When Apple introduced the original iPhone the world of smartphones was changed forever. It was arguably one of the most influential products ever to be launched in the consumer technology market and set the benchmark for the (many) smartphones that came after it. It sold millions worldwide and helped to make Apple into the world's most valuable company. Yet how Apple and its visionary then leader, Steve Jobs, managed the innovation process remained something of a secret for years after the product's launch. Originally visualised as a tablet computer, work on the iPhone started partly because of the success of the firm's earlier product, the iPod music player. Yet, a technological breakthrough (multi-touch display) allowed the company to change course. As Steve Jobs said later, 'I had this idea about having a glass display, a multi-touch display you could type on. I asked our people about it. And six months later they came back with this amazing display [when] we got inertial scrolling working and some other things, I thought, "my god, we can build a phone with this" and we put the tablet aside, and we went to work on the phone.' However, making the multi-touch display a working proposition was challenging for Apple's engineering team. They had to create an entirely new way in which users could interact with their phones. There were many novel unsolved problems to overcome. Every single part of the design had to be rethought to adapt to touch. For example, engineers had to make scrolling work on the iPhone not only when a user's finger moved up and down, but also when a user's thumb moved in an arc across the screen. In addition, there were many other obstacles to overcome, some which seemed almost insurmountable. Sir Jonathan Ive, senior vice-president of design at Apple, has admitted that issues with the touchscreen were so difficult that it brought the project to the brink of being aborted. 'There were multiple times when we nearly shelved it because there were fundamental problems that we couldn't solve,' said Sir Jonathan, 'I would put the phone to my ear and my ear dialled a number. The challenge is that you then have to detect all sorts of ear shapes, chin shapes, skin colour and hairdos. We had to develop technology, basically a number of sensors, to inform the phone that 'this is now going up to an ear, please deactivate the touchscreen'.

The aesthetics of the iPhone were treated as being just as important as the iPhone technology. This was the responsibility of Apple's secretive industrial design group. Apple designer, Christopher Stringer said that their objective was to create a 'new, original, and beautiful object [that was] so wonderful that you couldn't imagine how you'd follow it'. The design group, Stringer explained, comprised 16 'maniacal' individuals who shared one singular purpose – to 'imagine products that don't exist and guide them to life.' They worked closely together, often gathering around a 'kitchen table' where team members exchanged ideas, often in a brutally honest way. To the designers, even the tiniest of details were important. They often would create up to 50 designs of a single component before moving on to computer-aided design modelling and the creation of physical mock-ups.

The fact that the Apple designers overcame several technology and production bugs during its development is partly a testament to their design team's belief, both in their technological skills and in their understanding of what people will buy. Yet Apple avoids conducting market research when designing its products, a policy introduced by Steve Jobs, its late chief executive. '*We absolutely don't do focus groups,*' said Ive. '*That's designers and leaders abdicating responsibility. That's them looking for an insurance policy, so if something goes wrong, they can say, well this focus group says that only 30 per cent of people are offended by this and, look, 40 per cent think it's OK. What a focus group does is that it will guarantee mediocrity*.'

EXAMPLE

The sad tale of Kodak and its digital camera[3]

The once mighty Eastman Kodak Company dominated the photographic and film markets for decades, but no longer does so. Thirty years ago it employed over 140,000 people and made substantial profits; by 2010 it had shrunk to around 19,000, with regular quarterly losses. This dramatic fall from grace is usually put down to the company's failure to see the approach of digital photography or fully appreciate how it would totally undermine Kodak's traditional products. Yet, ironically, Kodak was more ahead of its competitors than most people outside the company realised. It actually invented the digital camera. Sadly, though, it lacked the foresight to make the most of it. For years the company had, as one insider put it, 'too much technology in its labs rather than in the market'.

It was in 1975 when a newly hired scientist at Kodak, Steve Sasson, was given the task of researching how to build a camera using a comparatively new type of electronic sensor – the charged-couple device (CCD). He found little previous research so he used the lens from a Kodak motion-picture camera, an analogue-to-digital convertor, some CCD chips and some digital circuitry that he made himself. By December 1975, he had an operational prototype. Yet the advance was largely, although not completely, ignored inside the company. '*Some people talked about reasons it would never happen, while others looked at it and realised it was important*,' he said. He also decided not to use the word "*digital*" to describe his trial product. '*I proposed it as filmless photography, an electronic stills camera. Calling it "digital" would not have been an advantage. Back then 'digital' was not a good term. It meant new, esoteric technology*'. Some resistance came from genuine, if mistaken, technical reservations. Others feared the magnitude of the changes that digital photography could bring. Objections. . .' *were coming from the gut: a realisation that [digital] would change everything – and threaten the company's entire film-based business model*.' Some see Kodak's reluctance to abandon its traditional product range as understandable. It was making vast profits and as late as 1999 it was making over three billion dollars from film sales. Todd Gustavson, curator of technology at the George Eastman House museum said that, '*Kodak was almost recession-proof until the rise of digital. A film-coating machine was like a device that printed money*.' So Kodak's first digital camera, the Quicktake, was licensed to and sold by Apple in 1994. In 2012 Kodak filed for bankruptcy protection.

What do these two examples have in common?

It may seem that the answer to the question of what the two examples have in common is 'nothing'. One is a highly successful company that revolutionised the design of the products into its markets, while the other had to protect itself from its creditors because it failed to innovate. Yet they do have something in common: both invented product technologies that were radically different from what had gone before. Apple developed an innovative set of smartphone technologies; Kodak invented the digital camera. Both of these innovations went on to change their industries. The real difference between the two examples lies not in whether they were innovative (both were) but in how they capitalised on their innovation. Apple sees technical innovation (and product design) as the main driver of its success. The knowledge gained in one new product forms the basis of innovations in other, different, products. It is willing to invest in ideas that it believes will become important in the future. Kodak, although it did fund new research, had a culture that was afraid of, and threatened by, new technology. At a strategic level, Kodak seemed to be more concerned with maintaining 'business as usual'. Of course, to be fair, Apple had relatively little to lose when it started. Kodak, on the other hand, had a vast business built on the technology that digital photography would replace.

The innovation S-curve

When new ideas are introduced in services, products or processes, they rarely have an impact that increases uniformly over time. Usually performance follows an S-shaped progress. So, in the early stages of the introduction of new ideas, although, often large, amounts of resources, time and effort are needed to introduce the idea, relatively small performance improvements are experienced. However, with time, as experience and knowledge about the new idea grow, performance increases. Nevertheless, as the idea becomes established, extending its performance further becomes increasingly difficult, see Figure 3.3 (a). When one idea reaches its mature, 'levelling off' period, it is vulnerable to a further new ideas being introduced that, in turn, moves through its own S-shaped progress. This is how innovation works; the limits of one idea being reached prompts a newer, better idea, with each new S-curve requiring some degree of re-design, see Figure 3.3 (b).

OPERATIONS PRINCIPLE

Usually the performance of each new innovation follows an S-shaped progress.

Incremental or radical innovation

An obvious difference between how the pattern of new ideas emerges in different operations or industries is the rate and scale of innovation. Some industries, such as telecommunications, enjoy frequent and often significant innovations. Others, such as house building do have innovations, but they are usually less dramatic. So some innovation is radical, resulting in

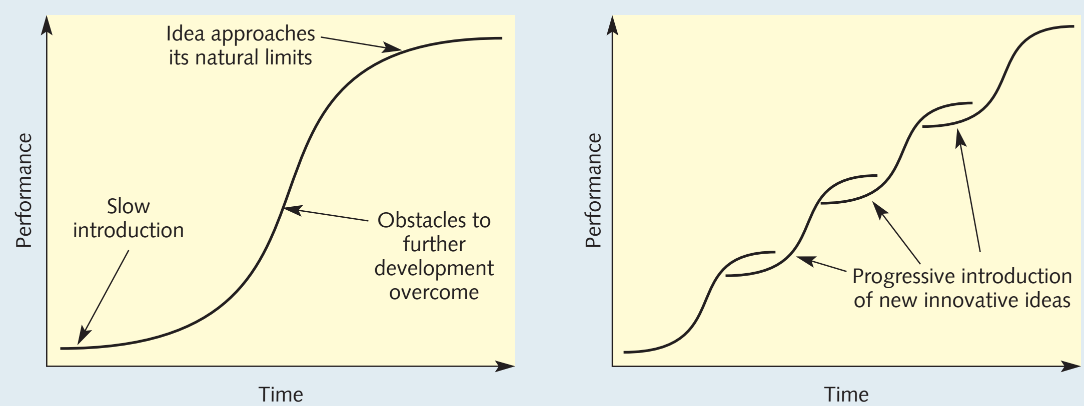

(a) The basic S-shaped improvement in performance

(b) Innovation following multiple S-shaped curves

Figure 3.3 The S-shaped curve of innovation

discontinuous, 'breakthrough' changes, while other innovations are more incremental leading to smaller, continuous changes. Radical innovation often includes large technological advancements that may require completely new knowledge and/or resources, making existing services and products obsolete and therefore non-competitive. Incremental innovation, by contrast, is more likely to involve relatively modest technological changes, build upon existing knowledge and/or resources, so existing services and products are not fundamentally changed. This is why established companies may favour incremental innovation because they have the experience to have built up a significant pool of knowledge (on which incremental innovation is based). In addition, established companies are more likely to have a mindset that emphasises continuity; perhaps not even recognising potential innovative opportunities (see the Kodak example earlier). New entrants to markets, however, have no established position to lose, nor do they have a vast pool of experience. They may be more likely to try for innovation that is more radical.

The Henderson–Clark Model

Although distinguishing between incremental and radical innovation is useful, it does not fully make clear why some companies succeed or fail at innovation. Two researchers, Henderson and Clark,[4] looked at the question of why some established companies sometimes fail to exploit seemingly obvious incremental innovations. They answered this question by dividing the technological knowledge required to develop new services and products into, 'knowledge of the components of knowledge' and 'knowledge of how the components of knowledge link together'. They called this latter knowledge, 'architectural knowledge'. Figure 3.4 shows what has become known as the Henderson–Clark Model. It refines the simpler idea of the split between incremental and radical innovation. In this model, incremental innovation is built upon existing component and architectural knowledge, whereas radical innovation changes both component and architectural knowledge. Modular innovation builds on existing architectural knowledge, but requires new knowledge for one or more components. By contrast, architectural innovation will have a great impact upon the linkage of components (or the architecture), but the knowledge of single components is unchanged.

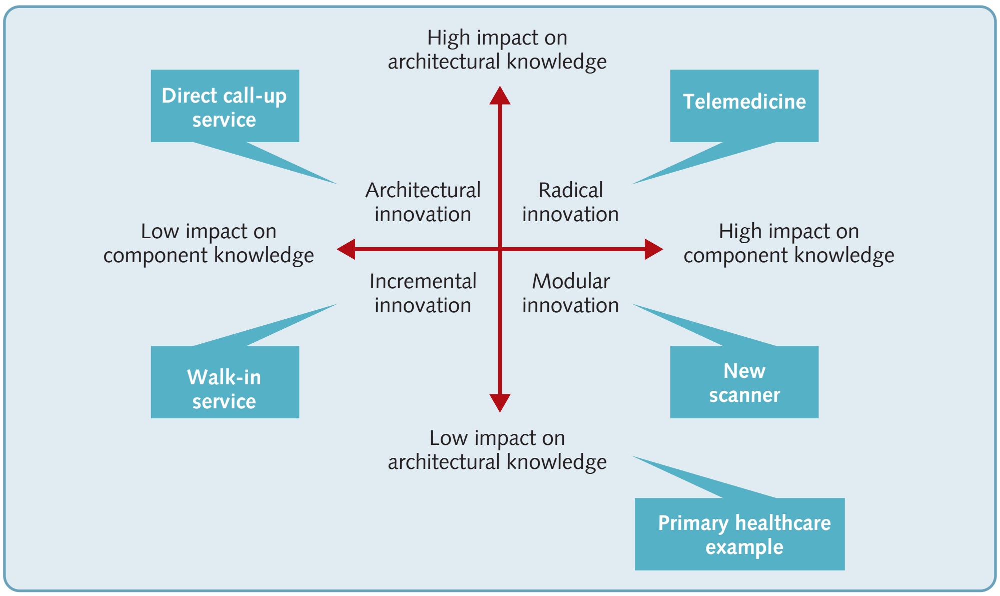

Figure 3.4 The Henderson–Clark Model

So, for example, in healthcare services, simple (but useful and possibly novel at the time) innovations in a primary care (general practitioner) doctors' clinic, such as on-line appointment websites would be classed as incremental innovation because neither any elements, nor the relationship between them is changed. If the practice invests in a new diagnostic heart scanner, that element of their diagnosis task has been changed and will probably need new knowledge, but the overall architecture of the service has not been changed. This innovation would be classed as 'modular'. An example of architectural innovation would be the practice providing 'walk-in' facilities in the local city centre. It would provide more or less the same service as the regular surgery (no new components), but the relationship between the service and patients has changed. Finally, if the practice adopted some of the 'telemedicine' technology, that would represent radical innovation.

OPERATIONS PRINCIPLE

Innovation can be classified on two dimensions, innovation of components of a design and innovation of the linkages between them.

Product and service design innovation as a process

Although organisations will have their own particular ways of managing innovation and design, the innovation process is essentially very similar across a whole range of industries. More importantly, it is a process that is seen as increasingly important. The aim of the innovation process is firstly to create offerings that exceed customers' expectations in terms of quality, speed, dependability, flexibility, cost and sustainability; and secondly, to ensure that competitors find these offerings hard to imitate, substitute, or gain access to. With increasingly demanding customers, higher levels of competition and shorter product-service lifecycles seen in many markets worldwide, organisations that can master the art of innovation will generate significant competitive advantage. In addition, if an organisation is able to consistently deliver innovations ahead of its competitors, it is more likely to set industry standards, which others are often forced to follow.

Like any other process, the innovation activity can and should be managed – after all, successful outcomes, in the form of service and product offerings, are significantly reliant on an effectively designed innovation process. While innovation projects clearly vary, there is sufficient commonality in them to be able to model the process and seek to improve it. Figure 3.5 shows the innovation activity as a process, with inputs and outputs, as in any other process. Inputs come both from within the organisation – employees, R&D, operations, marketing, human resources and finance; and from outside it – market research, customers, lead users, suppliers, competitors, collaborators and wider stakeholders. The outputs (or outcomes) of the innovation process include not just the details of the offering, but also an understanding of its value proposition.

OPERATIONS PRINCIPLE

The innovation activity is a process that can be managed using the same principles as other processes.

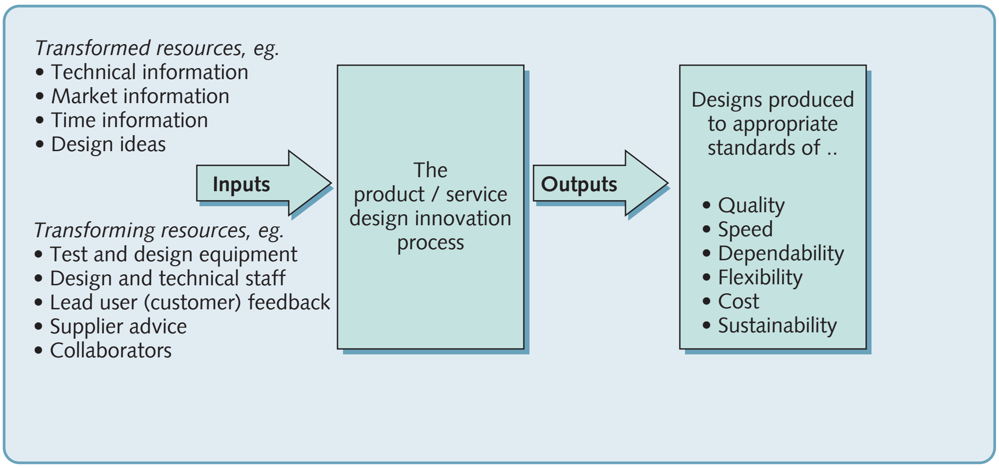

Figure 3.5 **The product and service design innovation activity as a process**

DIAGNOSTIC QUESTION

Are the product and service innovation process objectives specified?

OPERATIONS PRINCIPLE

Innovation processes can be judged in terms of their levels of quality, speed, dependability, flexibility, cost and sustainability.

The performance of the innovation process can be assessed in much the same way as we would consider the products and services that results from it, namely in terms of quality, speed, dependability, flexibility, cost, and sustainability. These performance objectives have just as much relevance for innovation as they do for the ongoing delivery of offerings once they are introduced to the market.

What is the quality of the innovation process?

Design quality is not always easy to define precisely, especially if customers are relatively satisfied with existing service and product offerings. Many software companies talk about the 'I don't know what I want, but I'll know when I see it' syndrome, meaning that only when customers use the software are they in a position to articulate what they do or don't require. Nevertheless, it is possible to distinguish high-and low-quality designs (although this is easier to do in hindsight) by judging them in terms of their ability to meet market requirements. In doing this, the distinction between the specification quality and the conformance quality of designs is important. No business would want a design process that was indifferent to 'errors' in its designs, yet some are more tolerant than others. For example, in pharmaceutical development the potential for harm is particularly high because drugs directly affect our health. This is why the authorities insist on such a prolonged and thorough design process. Although withdrawing a drug from the market is unusual, it does occasionally occur. Far more frequent are the 'product recalls' that are relatively common in, say, the automotive industry. Many of these are design related and the result of 'conformance' failures in the design process. The 'specification' quality of design is different. It means the degree of functionality, or experience, or aesthetics, or whatever the product or service is primarily competing on. Some businesses require product or service designs that are relatively basic (although free from errors), while others require designs that are clearly special in terms of the customer response they hope to elicit.

What is the speed of the innovation process?

The speed of innovation matters more to some industries than others. For example, innovation in construction and aerospace happen at a much slower pace than in clothing or microelectronics. However, rapid innovation or 'time-based competition' has become the norm for an increasing number of industries. Sometimes this is the result of fast-changing consumer fashion. Sometimes a rapidly changing technology base forces it. Telecoms, for example, are updated frequently because their underlying technology is constantly improving. Yet, no matter what the motivation, fast design brings a number of advantages:

- *Early market launch*: increasing the speed of services or product development allows earlier marketing of new offerings that may command price premiums and generate revenues for longer.
- *Starting design late*: alternatively starting the design process later may have advantages, especially where either the nature of customer demand or the availability of technology is

uncertain and dynamic, so fast design allows design decisions to be made closer to the time when service and product offerings are introduced to the market.

- *Frequent market stimulation*: rapid innovations allow frequent new or updated offerings to be introduced into the market.

What is the dependability of the innovation process?

Rapid innovation processes that cannot be relied on to deliver dependably are, in reality, not fast at all. Design schedule slippage can extend design times, but worse, a lack of dependability adds to the uncertainty surrounding the innovation process. Conversely, processes that are dependable minimise design uncertainty. Unexpected technical difficulties, such as suppliers who themselves do not deliver solutions on time, customers or markets that change during the innovation process itself, and so on, all contribute to an uncertain and ambiguous design environment. Professional project management (see Chapter 15) of the innovation process can help to reduce uncertainty and prevent (or give early warning of) missed deadlines, process bottlenecks and resource shortages. However, external disturbances to the innovation process will remain. These may be minimised through close liaison with suppliers and market or environmental monitoring. Nevertheless, unexpected disruptions will always occur and the more innovative the design, the more likely they are to occur. This is why flexibility within the innovation process is one of the most important ways in which dependable delivery of new service and product offerings can be ensured.

What is the flexibility of the innovation process?

Flexibility in the innovation process is the ability to cope with external or internal change. The most common reason for external change is that markets, or specific customers, change their requirements. Although there may be no need for flexibility in relatively predictable markets, it is clearly valuable in more fast-moving and volatile markets, where one's own customers and markets change, or where the designs of competitors' offerings dictate a matching or leapfrogging move. Internal changes include the emergence of superior technical solutions. In addition, the increasing complexity and interconnectedness of service and product components in an offering may require flexibility. A bank, for example, may bundle together a number of separate services for one particular segment of its market. Privileged account holders may obtain special deposit rates, premium credit cards, insurance offers, travel facilities, and so on, together in the same package. Changing one aspect of this package may require changes in other elements. So extending the credit card benefits to include extra travel insurance may also mean the redesign of the separate insurance element of the package. One way of measuring innovation flexibility is to compare the cost of modifying a design in response to such changes against the consequences to profitability if no changes are made. The lower the cost of modifying an offering in response to a given change, the higher is the level of flexibility.

What is the cost of the innovation process?

The cost of innovation is usually analysed in a similar way to the ongoing cost of delivering offerings to customers. These cost factors are split up into three categories: the cost of buying the inputs to the process; the cost of providing the labour in the process; and the other general overhead costs of running the process. In most in-house innovation processes, the last two costs outweigh the first.

One way of thinking about the effect of the other innovation performance objectives on cost is shown in Figure 3.6. Whether caused by quality errors, an intrinsically slow innovation process, a lack of project dependability, or delays caused through inflexibility, the end result is

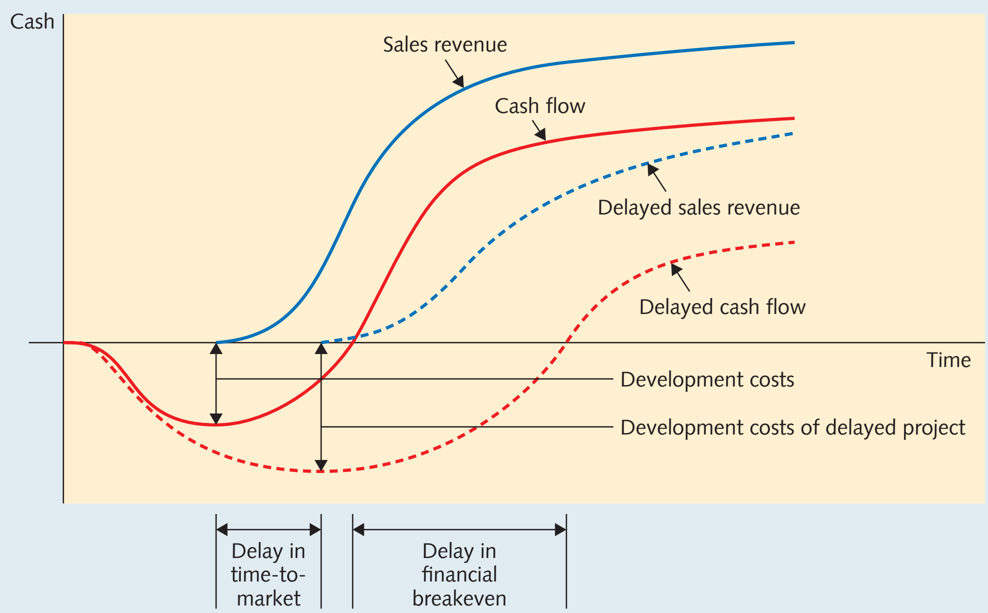

Figure 3.6 **Delay in time to market of new innovations not only reduces and delays revenues; it also increases the costs of development. The combination of both of these effects usually delays the financial break-even point far more than the delay in the launch**

that the design is late. Delayed completion of the design results in both more expenditure on the design and delayed (and probably reduced) revenue. The combination of these effects usually means that the financial break-even point for a new offering is delayed far more than the original delay in its launch.

What is the sustainability of the innovation process?

The product and service innovation process is particularly important in ultimately impacting on the ethical, environmental and economic well-being of stakeholders. Some innovation activity particularly focuses on the ethical dimension of sustainability. Banks have moved to offer ethical investments that seek to maximise social benefit as well as financial returns (avoiding businesses involved in weaponry, gambling, alcohol and tobacco, for example). Other examples of ethically focused innovations include the development of 'fair-trade' products such as bananas, tea, coffee, flowers, chocolate, cotton, and handicrafts; clothing manufacturers establishing ethical trading initiatives with suppliers; supermarkets ensuring animal welfare for meat and dairy, and paying fair prices for vegetables; and online companies establishing customer complaint charters. Innovation may also focus on changing materials in the design to reduce its environmental burden. Examples include the use of organic cotton or bamboo in clothing; wood or paper from managed forests used in garden furniture, stationery and flooring; recycled materials for carrier bags; and natural dyes in clothing, curtains and upholstery.

Other innovations may focus more on the use stage of an offering. The MacBook Air, for example introduced a power management system that reduced the power requirements. In the detergent industry, Persil and Ariel have developed products that allow clothes to be washed at lower temperatures. Architects are increasingly designing houses that can operate with minimal energy, or use sustainable sources of energy such as solar panels. Some innovations focus on making product components within an offering easier to recycle or remanufacture once they have reached the end of their life. Some food packaging is designed to break down easily when disposed of, allowing its conversion into high-quality compost. Mobile phones can be designed to be taken apart at the end of their life, so valuable raw materials can be reused in new phones. In the car industry, over 75 per cent of materials are now recycled.

EXAMPLE

Product innovation for the circular economy[5]

Design innovation is not just confined to the initial conception of a product; it also applies to the end of its life. This idea is often called 'designing for the circular economy'. The 'circular economy' is proposed as an alternative to the traditional linear economy (or make-use-dispose as it is termed). The idea is to keep products in use for as long as possible, extract the maximum value from them while in use, and then recover and regenerate products and materials at the end of their service life. The circular economy is, however, much more than a concern for recycling as opposed to disposal. The circular economy examines what can be done right along the supply and use chain so that as few resources as possible are used, then (and this is the important bit) recover and regenerate products at the end of their conventional life. This means designing products for longevity, reparability, ease of dismantling and recycling.

Typical of the companies that have adopted this idea is Newlife Paints, on the south coast of England; it 'remanufactures' waste water-based paint back into a premium-grade emulsion. All products in the company's paint range guarantee a minimum of 50 per cent recycled content, made up from waste paint diverted from landfill or incineration. Industrial Chemist, Keith Harrison, started the company. His wife encouraged him to clean up his unruly garage after many years of do-it-yourself projects. After realising that the stacked-up tins of paint represented a shocking waste, he began to search for a sensible and environmentally responsible solution to waste paint. '*I kept thinking I could do something with it, the paint had an intrinsic value. It would have been a huge waste just to throw it away*.' After two years of research, Keith successfully developed his technology, which involves removing leftover paint from tins that have been diverted from landfill, and blending and filtering them to produce colour-matched new paints. The company has also launched a premium brand, aimed at affluent customers with a green conscience, called Reborn Paints; a development that was partly funded by Akzo Nobel, maker of Dulux Paints. Keith now licenses his technology to companies including the giant waste company Veolia. '*By licensing we can have more impact and spread internationally*,' he said. He also points out that manufacturers could plan more imaginatively for the afterlife of their products. For example, simply adding more symbols to packs to assist sorting waste paints into types would help.

DIAGNOSTIC QUESTION

Is the product and service innovation process defined?

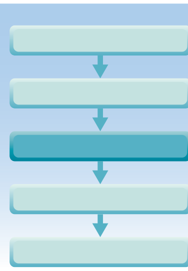

To create a fully specified service or product offering, potential designs must pass through several stages. A typical innovation process is shown in Figure 3.7. Although not all companies use these exact stages and there are often rework loops involved, there is still considerable similarity between the stages and sequence of the innovation process. Furthermore, they all

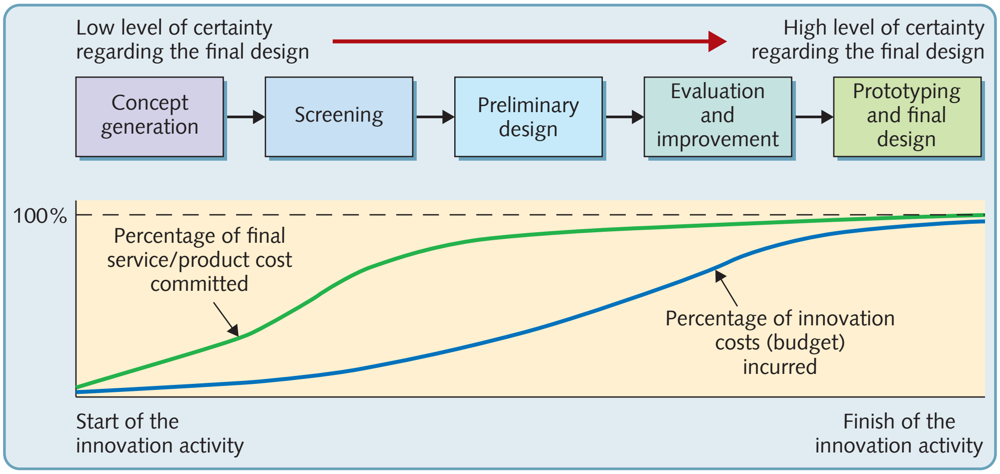

Figure 3.7 The stages in a typical innovation process and the design funnel effect – progressively reducing the number of possibilities until the final design is reached

OPERATIONS PRINCIPLE

Innovation processes involve a number of stages that move an innovation from a concept to a fully specified state.

have the same underlying principle; that over time an original idea, or 'concept', is refined and made progressively more detailed until it contains sufficient information for turning into an actual service, product, or process. At each stage in this progression the level of certainty regarding the final design increases as design options are discarded. The final design will not be evident until the very end of the process. Yet, many of the decisions that affect the eventual cost of delivery are made relatively early. For example, choosing to make a mobile telephone's case out of a magnesium alloy will be a relatively early decision that may take little investigation. Yet, although accounting for a small part of the total design budget, this decision may go a long way to determining the final cost of the phone. The difference between the 'budget spend' of the innovation process and the actual costs committed by the innovation process are also shown in Figure 3.7.

Concept generation

Concept generation is all about ideas and ideas can come from anywhere. Often the expectations within organisations are that ideas will emerge from the research and development (R&D) or market research departments. However, this ignores the huge potential of other internal sources of innovation. Front-line service providers, in particular, are able to provide deep insights into what customers require based on informal interactions. Similarly, while many customer complaints are dealt with at a relatively operational level, they have the potential to act as a useful source of customer opinion within the innovation process. Suppliers can also be valuable in the innovation process because of their potential to improve the quality of products and services, minimise time to market and spread the cost and risks of innovation.

'Lead users' and 'harbingers of failure'

Although the marketing function is generally responsible for identifying new service or product opportunities, often by using formal structured market research tools, ideas may come less formally; for example, listening to customers on a day-to-day basis from every-day transactions or from complaints. A particularly useful source of customer-inspired innovation, especially for products and services subject to rapid change, are so-called 'lead users'.[6] 'Lead users' are users who are ahead of the majority of the market on a major market trend, and who also have a high incentive to innovate.[7] Producers seeking user innovations to manufacture try to source innovations from lead users – because these will be most profitable to manufacture. These customers, unlike most customers, have the real-world experience needed to problem solve and provide accurate data to inquiring market researchers. Since these lead users will be familiar both with the positives and negatives of the early versions of products and services, they are a particularly valuable source of potential innovative ideas.

If lead users can bring insights because of their expertise, by contrast another category of customer may be valuable because of their ability consistently to make bad purchase decisions. These customers have been termed 'harbingers of failure'. One study[8] claims that the same group of consumers has a tendency to purchase all kinds of failed products, time after time, flop after flop. As one of the authors of the study put it, 'These harbingers of failure have the unusual property that they keep on buying products that are taken from the shelves. These star-crossed consumers can sniff out flop-worthy products of all kinds. If you're the kind of person who bought something that really didn't resonate with the market, say, coffee-flavoured Coca-Cola, then that also means you're more likely to buy a type of toothpaste or laundry detergent that fails to resonate with the market.'

Ideas management

Obtaining new product or service ideas (or indeed any innovative ideas) from employees can also be a rich source of innovation. For example, the 3M Corporation has been highly successful in generating innovations by introducing formal incentives to encourage employee engagement. Employee-sourced ideas were traditionally done through paper-based 'suggestion schemes' where employees placed their ideas in a 'suggestion box'. Such schemes were often only partly effective, yielding few, low-quality ideas. Unless the running of the scheme was well-resourced, it could be difficult to guarantee that all ideas were evaluated consistently and quickly. The scheme could lose credibility unless employees could track their ideas to confirm that they 'didn't just disappear'. However, the advent of 'idea management' software tools has overcome some of these difficulties. Ideas management systems are a type of enterprise software (often web-based) that can help operations to collect ideas from employees, assess them and, if appropriate, implement them quickly and efficiently. Such systems can track ideas all the way, through from inception to implementation, making it much easier to understand important performance measures such as where ideas are being generated, how many ideas submitted are actually implemented, the estimated cost savings from submitted ideas and any new revenues generated by implemented ideas. Often ideas management systems focus ideas on specific organisational targets and objectives, which it is claimed improve both the quality and quantity of ideas, when compared with 'open' suggestion schemes.

Concept screening

Concept screening is the first stage of implementation where potential innovations are considered for further development. It is not possible to translate all concepts into viable product-service packages. Organisations need to be selective. For example, DuPont estimates that the ratio of concepts to marketable offerings is around 250-to-1. In the pharmaceuticals industry, the ratio is closer to 10,000-to-1. The purpose of concept-screening is to take initial concepts and evaluate them for their feasibility (can we do it?), acceptability (do we want to do it?) and vulnerability (what are the risks of doing it?). Concepts may have to pass through many different screens, and several functions might be involved. Table 3.1 gives typical feasibility, acceptability and vulnerability questions for marketing, operations and finance functions.

During concept screening a key issue to consider is deciding how big the innovation should be and where it should focus – innovation to the customer offering as opposed to innovation

Table 3.1 Some typical evaluation questions for marketing, operations and finance

Evaluation criteria	*Marketing*	*Operations*	*Finance*
Feasibility	Is the market likely to be big enough?	Do we have the capabilities to deliver it?	Do we have access to finance to develop and launch it?
Acceptability	How much market share could it gain?	How much will we have to reorganise our activities to deliver it?	How much financial return will there be on our investment?
Vulnerability	What is the risk of it failing in the marketplace?	What is the risk of us being unable to deliver it acceptably?	How much money could we lose if things do not go to plan?

to the process of delivery. The vast majority of innovation is continuous or incremental in nature. Here the emphasis is on steady improvement to existing offerings and to the processes that deliver them. This kind of approach to innovation is very much reflected in the lean and total quality management perspectives. On the other hand, some innovation is discontinuous and involves radical change that is 'new to the world'. Discontinuous innovation is relatively rare – perhaps 5–10 per cent of all innovations could be classified as such – but creates major challenges for existing players within a market. This is because organisations are often unwilling to disrupt current modes of working in the face of a barely emerging market, but by the time the threat has emerged more fully it may be too late to respond. Clayton Christensen refers to this problem as the Innovator's Dilemma, which supports renowned economist Joseph Schumpeter's ideas that innovation should be a process of 'creative destruction'.

Preliminary design

Having generated one or more appropriate concepts the next stage is to create preliminary designs. For service-dominant offerings, this may involve documentation in the form of job instructions or 'service blueprints'. For product-dominant offerings, preliminary design involves defining product specifications (McDonalds has over 50 specifications for the potatoes used for its fries) and the bill of materials, which details all the components needed for a single product. At this stage, there are significant opportunities to reduce cost through design simplification. The best innovations are often the simplest. Designers can adopt a number of approaches to reduce design complexity. These include standardisation, commonality, modularisation and mass customisation.

OPERATIONS PRINCIPLE

A key innovation objective should be the simplification of the design through standardisation, commonality, modularisation and mass customisation.

Standardisation

This is an attempt to overcome the cost of high variety by standardising offerings, usually by restricting variety. Examples include fast-food restaurants; discount supermarkets or telephone-based insurance companies. Similarly, although everybody's body shape is different, garment manufacturers produce clothes in a limited number of sizes. The range of sizes is chosen to give a reasonable fit for most, but not all, body shapes. Controlling variety is an important issue for most businesses, which all face the danger of allowing variety to grow excessively. Many organisations have significantly improved their profitability by careful variety reduction, often by assessing the real profit or contribution of each service or product.

Commonality

Common elements are used to simplify design complexity. If different services and products can draw on common components, the easier it is to deliver them. An example of this is Airbus, the

European aircraft maker, which designed its aircraft with a high degree of commonality using fly-by-wire technology. This meant that ten aircraft models featured virtually identical flight decks, common systems and similar handling characteristics. The advantages of commonality for the airline operators include a much shorter training time for pilots and engineers when they move from one aircraft to another. This offers pilots the possibility of flying a wide range of routes from short-haul to ultra-long haul and leads to greater efficiencies because common maintenance procedures can be designed with maintenance teams capable of servicing any aircraft in the same family. In addition, when up to 90 per cent of all parts are common within a range of aircraft, there is a reduced need to carry a wide range of spare parts. Similarly, Hewlett-Packard, and Black and Decker have used common platforms to reduce innovation costs.

Modularisation

This is a method of balancing two opposite forces: standardisation and customisation. It involves designing standardised 'sub-components' of an offering that can be put together in different ways. For example, the package holiday industry can assemble holidays to meet a specific customer requirement, from pre-designed and purchased air travel, accommodation, insurance, and so on. Similarly, in education modular courses are increasingly used that allow 'customers' choice but permit each module to have economical volumes of students. Dell, a pioneer in computer manufacture, used the same logic for products, drawing together interchangeable sub-assemblies, manufactured in high volumes (and therefore lower cost), in a wide variety of combinations. In a similar way software engineering often involves modularisation to bring some degree of order to the development of large and complex pieces of software. These can often involve a large number of programmers. Here modularisation allows the sometimes thousands of lines of code to be broken up and organised by the task it performs.

Mass customisation[9]

Flexibility in design can allow the ability to offer different things to different customers. Normally, high variety means high cost, but some companies have developed their flexibility in such a way that customised offerings are produced using high-volume processes and thus costs are minimised. This approach is called mass customisation. For example, Paris Miki, an up-market eyewear retailer that has the largest number of eyewear stores in the world, uses its own 'Mikissimes Design System' to capture a digital image of the customer and analyse facial characteristics. Together with a list of customers' personal preferences, the system then recommends a particular design and displays it on the image of the customer's face. In consultation with the optician, the customer can adjust shapes and sizes until choosing the final design. The frames are assembled within the store from a range of pre-manufactured components and the lenses ground and fitted to the frames. The whole process takes around an hour.

EXAMPLE

Customising for kids

A major challenge facing global programme makers is achieving economies of scale that result from high-volume production, while allowing customisation of programmes to suit different markets. 'Art Attack!', made for the Disney Channel – a children's TV channel shown around the world – used the concept of mass customisation to meet this challenge. Typically, over two hundred episodes of the show are made in six different language versions. About 60 per cent of each show is common across all versions. Scenes that do not include anyone speaking or where the presenter's face is not visible are shot separately. For example, in an episode showing how to make a simple cardboard model , all versions will share the scenes where the presenter's hands only are visible. Commentary in the appropriate language is over-dubbed onto the scenes that are edited seamlessly with other shots of the appropriate presenter. The final product will have the head and shoulders of Brazilian, French, Italian, German, or Spanish presenters flawlessly mixed with the same pair of hands constructing the model. Local viewers in each market saw what appeared to be a highly customised show, yet the cost of making each episode was about one third of producing truly separate programmes for each market.

Design evaluation and improvement

The purpose of this stage in the design activity is to take the preliminary design and see if it can be improved before the service-product offering is tested in the market. A number of techniques can be employed at this stage to evaluate and improve the preliminary design. Perhaps the best known is quality function deployment (QFD). The key purpose of QFD is to try to ensure that the eventual innovation actually meets the needs of its customers. It is a technique that was developed in Japan at Mitsubishi's Kobe shipyard and used extensively by Toyota, the motor vehicle manufacturer, and its suppliers. QFD is also known as the 'house of quality' (because of its shape) and the 'voice of the customer' (because of its purpose). The technique tries to capture what the customer needs and how it might be achieved. Figure 3.8 shows a simple QFD matrix used in the design of a promotional USB data storage pen. It is a formal articulation of how designers see the relationship between the requirements of the customer and the design characteristics of the offering.

At this stage in the process, both creativity and persistence are needed to move from a potentially good idea to a workable design. One product has commemorated the persistence of its design engineers in its company name. In 1953 the Rocket Chemical Company set out to create a rust-prevention solvent and degreaser to be used in the aerospace industry. Working in their lab in San Diego, California, they made 40 attempts to work out the water displacing formula. So they called the product 'WD-40®', which literally stands for Water Displacement,

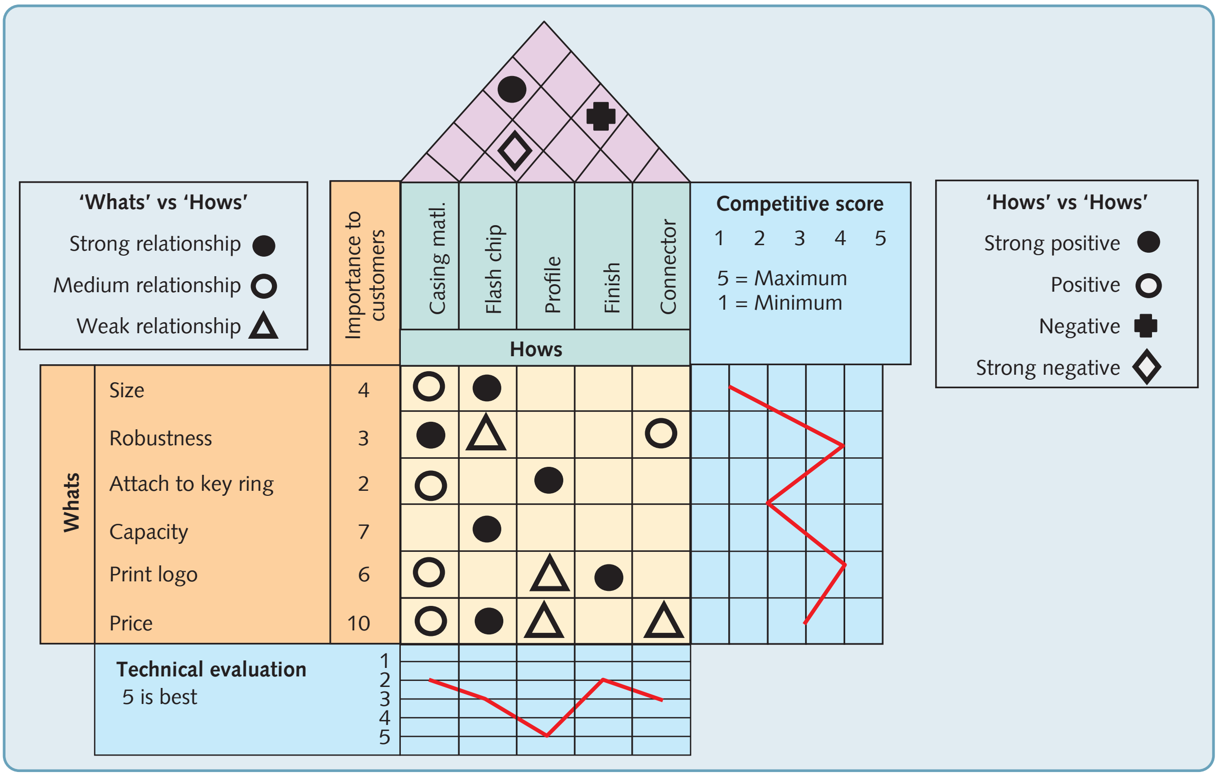

Figure 3.8 A QFD matrix for a promotional USB data storage pen

fortieth attempt. Originally used to protect the outer skin of the Atlas Missile from rust and corrosion, the product worked so well that employees kept taking cans home to use for domestic purposes. Soon after, the product was launched with great success into the consumer market.

In fact, it's not just persistence that is important in the innovation process – failure itself may be beneficial if organisations can spot potential. Sometimes, when a design fails, it represents an opportunity to re-think the concept itself. For example, Pritt Stick, the world's first glue stick, was originally intended to be a super-glue, but product testing proved unsatisfactory. Henkel changed the product concept and successfully marketed the product as 'the non-sticky sticky stuff'. Similarly, a group of chemists working for the pharmaceutical giant Pfizer developed a new drug called 'Sildenafil'. Originally intended to help individuals with hypertension (high blood pressure) and angina, clinical trials of the drug proved unsuccessful, though doctors noticed a side effect of penile erection. Seeing the potential of this 'failed' innovation process, Pfizer marketed the drug, as Viagra, for erectile dysfunction. In just two years, sales of Viagra had topped $1billion and the product dominated the market.

Prototyping and final design

At around this stage in the innovation activity it is necessary to turn the improved design into a prototype in order to test it. Product prototypes include everything from clay models to computer simulations. For more service-dominant offerings, prototyping often involves the actual implementation of the service on a pilot basis. For example, a retailer may organise piloting new services packages in a small number of stores in order to test customers' reaction to them. It may also be possible to 'virtually' prototype in much the same way as a product. This is a familiar idea in some industries such as magazine publishing, where images and text can be rearranged and subjected to scrutiny prior to existing in any physical form, allowing amendments to be made right up to the point of production. Although the means of prototyping may vary, the principle is always the same; do whatever one can to test out the innovation prior to delivery. Doing so minimises the likelihood of launching an unsuccessful innovation into the market and the financial and reputational risks that come with such a scenario.

Alpha and beta testing

A distinction that originated in the software development industry, but has spread into other areas is that between the alpha and beta testing of a product or service. Most software products include both alpha and beta test phases, both of which are intended to uncover 'bugs' (errors) in the product. Not surprisingly, alpha testing comes before beta testing. Alpha testing is essentially an internal process where the developers or manufacturers (or sometimes an outside agency they have commissioned) examine the product for errors. Generally, it is also a private process, not open to the market or potential customers. Although it is intended to look for errors that otherwise would emerge when the product is in use, it is, in effect, performed in a virtual or simulated environment, rather than in 'the real world'. After alpha testing, the product is released for beta testing; that is, when the product is released for testing by selected customers. It is an external 'pilot-test' that takes place in 'real world' (or near real world, because it is still a relatively small, and short, sample) before commercial production. By the time a product gets to the beta stage most of the worst defects should have been removed, but the product may still have some minor problems that may only become evident with user participation. This is why beta testing is almost always performed at the user's premises without any of the development team present. Beta testing is sometimes called 'field testing', pre-release testing, customer validation, customer acceptance testing, or user acceptance testing.

DIAGNOSTIC QUESTION

Are the resources for developing product and service innovation adequate?

For any process to operate effectively, it must be appropriately designed and resourced. Innovation processes are no different. (The detailed principles of process design are discussed in Chapters 5 and 6.) However, because innovation processes are often an operation within the business in their own right, there are some more strategic resourcing issues to consider – how much capacity to devote to innovation, how much of the innovation activity to outsource and what kinds of technology to use in the innovation process.

> OPERATIONS PRINCIPLE
>
> *For innovation processes to be effective they must be adequately resourced.*

Is there sufficient innovation capacity?

In general, capacity management involves deciding on the appropriate level of capacity needed by a process and how it can be adjusted to respond to changes in demand. For innovation, demand is the number of new designs needed by the business. The chief difficulty is that, even in very large companies, the rate of innovation is not constant. This means that product and service design processes are subjected to uneven internal 'demand' for designs, possibly with several new offerings being introduced to the market close together, while at other times little innovation is needed. This poses a resourcing problem because the capacity of an innovation activity is often difficult to flex. The expertise necessary for innovation is embedded within designers, technologists, market analysts, and so on. It may be possible to hire in some expertise as and when it is needed but much design resource is, in effect, fixed.

Such a combination of varying demand and relatively fixed design capacity means some organisations are reluctant to invest in innovation processes because they see it as an under-utilised resource. A vicious cycle can develop, in which companies fail to invest in innovation resources simply because many skilled design staff cannot be hired in the short-term, resulting in innovation projects overrunning or failing to deliver appropriate solutions. This, in turn, may lead to the company losing business or otherwise suffering in the marketplace, which makes the company even less willing to invest in innovation resources.

Should all innovation activities be done in-house?

Just as there are supply networks that produce services and products, there are also supply networks of knowledge that connect suppliers and customers in the innovation process; this is sometimes called the 'design (or development) network'. Innovation processes can adopt any position on a continuum of varying degrees of design engagement with suppliers, from retaining all the innovation capabilities in-house, to outsourcing all its innovation work. Between these extremes are varying degrees of internal and external capability. Figure 3.9 shows some of the more important factors that will vary depending on where an innovation process is on the continuum. Resources will be easy to control if they are kept in-house because they are closely aligned with the company's normal organisational structures, but control should be relatively loose because of the extra trust present in working with familiar colleagues. Outsourced design often involves greater control, with the use of contract with penalty clause for delay commonly used.

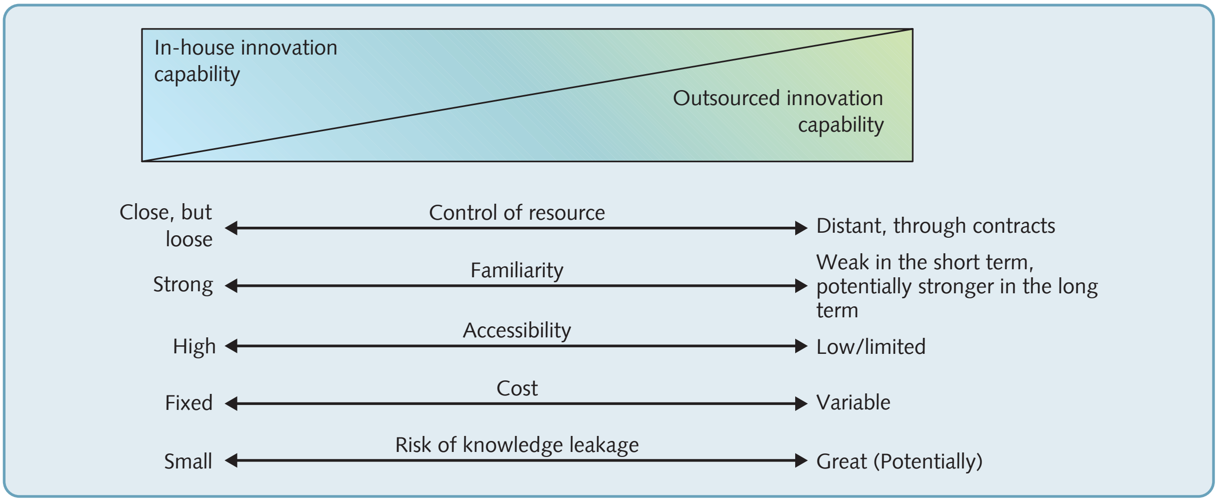

Figure 3.9 Some implications of the in-house–outsourced continuum

The overall cost of in-house versus outsourced innovation will vary, depending on the firm and the project. An important difference, however, is that external innovation tends to be regarded as a variable cost. The more external resources are used, the more these variable costs will be. In-house innovation is more of a fixed cost. Indeed a shift to outsourcing may occur because fixed design costs are viewed as too great. From an open innovation perspective, it is argued that firms should be willing to buy-in, or license, inventions from other organisations, rather than relying solely on innovations generated internally. Similarly, it may be beneficial to give access to under-used proprietary innovations through joint ventures, licensing, or spin-offs. However, a major inhibitor to open innovation is the fear of knowledge leakage. Firms become concerned that experience gained through collaboration with a supplier of design expertise may be transferred to competitors. There is a paradox here. Businesses usually outsource design primarily because of the supplier's capabilities that are themselves an accumulation of specialist knowledge from working with a variety of customers. Without such knowledge 'leakage', the benefits of the supplier's accumulated innovation capabilities would not even exist.

EXAMPLE

The most valuable patent in history

Alexander Graham Bell is widely credited with the invention of the first usable telephone in 1876. However, what is less widely known is that the Western Union Telegraph Company had the opportunity to buy into this innovation but chose not to – '*After careful consideration of your invention, which is a very interesting novelty, we have come to the conclusion that it has no commercial possibilities . . . we see no future for an electrical toy. . .* ' Within four years, 50,000 telephones had been sold in the USA and within twenty years the figure had reached 5 million. The patent, 174–465, is recognised as the most valuable in history and essentially created the telephone industry that we know today. This example illustrates the importance of drawing innovation from outside of the organisation. Perhaps one of the reasons why Western Union were so reluctant to invest in Bell's innovation was that is wasn't their idea! This so-called 'not-invented-here' syndrome is surprisingly common.

Alexander Graham Bell

Is appropriate technology being used in the innovation process?

Technology has become increasingly important in innovation activities. Simulation software, for example, is now common in the design of everything from transportation services through to chemical factories. These allow developers to make design decisions in advance of the creation of the actual product or service. They allow designers to work through the experience of using the service or product and learn more about how it might operate in practice. They can explore possibilities, gain insights and, most important, they can explore the consequences of their decisions. Innovation technologies are particularly useful when the design task is highly complex, because they allow developers to reduce their own uncertainty of how services or products will work in practice. Technologies also consolidate information on what is happening in the innovation process, thus presenting a more comprehensive vision within the organisation.

Computer-Aided Design (CAD)

The best-known innovation technology is computer-aided design (CAD). CAD systems store and categorise component information and allow designs to be built up on screen, often performing basic engineering calculations to test the appropriateness of proposed design solutions. They provide the computer-aided ability to create a modified product drawing and allow for the swift addition of conventionally used shapes to the computer-based representation of a product. Designs created on screen can be saved and retrieved for later use, which enables a library of standardised parts and components to be built up. Not only can this dramatically increase the productivity of the innovation process, it also aids the standardisation of parts in the design activity. For complex projects, the absolute size and interrelatedness of different aspects of the work requires sophisticated CAD systems to be successful.

Design for manufacture and assembly (DFMA) software is an extension of CAD that allows those involved in the innovation process to integrate their designs prior to manufacture. For example, DFMA could be used to see how a new gun turret design will fit onto armoured vehicles during production, before either the turret or the vehicle has been made physically. 3-D object modelling is a rapid prototyping technique aimed at reducing the time taken to create physical models of products. Designs from CAD are created using machines that build models by layering up extremely thin (usually 0.05mm) layers of photopolymer resin. An alternative to 3-D modelling is the use of virtual reality technologies. Here CAD information is converted into virtual images that can be viewed using 3-D glasses. This form of technology is more interactive that traditional CAD, as designers (or customers) can 'walk around' a design and get a better sense of what it looks and feels like to be in it. This means virtual reality technology is especially useful for designs that customers would be inside, such as sports venues, aeroplanes, buildings and amusement parks, for example.

Digital twins

Michael Grieves of Florida Institute of Technology's Centre for Lifecycle and Innovation Management popularised the term 'digital twin'. It is the combination of data and intelligence that represent the structure, context and behaviour of a physical system of any type, offering an interface that allows one to understand past and present operation, and make predictions about the future.[10] In other words, they are powerful digital 'replicas' that can be used instead of the physical reality of a product. Using the digital twin rather than the real product can significantly improve its operational performance without the expense of working on the real thing.[11] Not only that, but the use of the digital twin can extend throughout the product's life to provide valuable information to its user and evidence on how it is actually performing. Thus products live two parallel lives: one in real life, one in digital form. For example, digital twins could monitor and simulate possible future scenarios and predict the need for repairs and other problems before they occur. This allows design improvements to be made before a product is used by customers, as well as during its life.

Knowledge management technologies

In many professional service firms, such as management consultancies, design involves the evaluation of concepts and frameworks that can be used in client organisations to diagnose problems, analyse performance and construct possible solutions. They may include ideas of industry best practice, benchmarks of performance within an industry, and ideas that can be transported across industry boundaries. However, management consulting firms are often geographically dispersed and staff are often away from the office, spending most of their time in client organisations. This creates a risk for such companies of 'reinventing the wheel' continually. Most consultancy companies attempt to tackle this risk by using knowledge management routines based on their intranet capabilities. This allows consultants to put their experience into a common pool, contact other staff within the company who have skills relevant to a current assignment and identify previous similar assignments. In this way, information is integrated into the ongoing knowledge innovation process within the company and can be tapped by those charged with developing innovations.

DIAGNOSTIC QUESTION

Is the development of products and services and of the process that created them simultaneous?

Every product and service innovation has, eventually, to be created; this is why it is a mistake to separate the design of services and products from the design of the processes that will deliver them. Merging these two processes is sometimes called simultaneous (or interactive) design. Its key benefit is to reduce the elapsed time taken for the whole innovation activity. As noted earlier, reducing time to market (TTM) can give an important competitive advantage. If it takes a business two years to develop an offering with a given set of resources, it can only introduce new offerings every two years. If its rival can develop offerings in one year, it can introduce innovations into the market twice as often. This means the rival does not have to make such radical improvements in performance each time it introduces a new offering because it is introducing them more frequently. The factors that can significantly reduce time to market for innovations include the following:

- Integrating the design of the product-service offering and the design of the process used to create and deliver them.
- Overlapping the stages in the innovation process.
- An early deployment of strategic decision-making and resolution of design conflict.
- An organisational structure that reflects the nature of the offering.

OPERATIONS PRINCIPLE
Effective simultaneous innovation reduces time to market.

Integrating the design of the offering and design of the process

What looks good as an elegant offering on paper may prove difficult to create and deliver on an ongoing basis. Conversely, a process designed for one set of services or products may be incapable of creating different ones. It clearly makes sense to design offerings and operations processes together. For services, organisations have little choice but to do this because the process of delivery is usually part of the offering. However, it is useful to integrate the design of the offering and the process, regardless of the kind of organisation. The fact that many businesses do not do this is only partly because of their ignorance or incompetence. There are real barriers to doing it. First, the time scales involved can be very different. Offerings may be modified, or even redesigned, relatively frequently. The processes used to create and deliver an offering may be far too expensive to modify every time the offering changes. Second, the people involved with the innovation on the one hand, and ongoing process design on the other, are likely to be organisationally separate. Finally, it is sometimes not possible to design an ongoing process for the creation and delivery of services and products until they are fully defined.

Yet none of these barriers is insurmountable. Although it may not be possible to change ongoing processes every time there is a change to the offering, they can be designed to cope with a range of potential services and products. The fact that design staff and operations staff are often organisationally separate can also be overcome. Even if it is not sensible to merge the two functions, there are communication and organisational mechanisms to encourage the two functions to work together. Even the claim that ongoing processes cannot be designed until they know the nature of the offering is not entirely true. There can be sufficient clues emerging from innovation activities for process design staff to consider how they might modify ongoing processes. This is a fundamental principle of simultaneous design considered next.

Overlapping the stages of the innovation process

We have described the innovation process as a set of individual, predetermined stages, with one stage being completed before the next one commences. This step-by-step, or sequential, approach has been commonly applied in many organisations. It has some advantages. It is easy to manage and control innovation processes organised in this way because each stage is clearly defined. In addition, as each stage is completed before the next stage is begun, each stage can focus its skills and expertise on a limited set of tasks. The main problem of the sequential approach is that it is both time-consuming and costly. When each stage is separate, with a clearly defined set of tasks, any difficulties encountered during the design at one stage might necessitate the design being halted while responsibility moves back to the previous stage. This sequential approach is shown in Figure 3.10(a).

Often there is really little need to wait until the absolute finalisation of one stage before starting the next.[12] For example, perhaps while generating the concept, the evaluation activity of screening and selection could be started. It is likely that some concepts could be judged as 'non-starters' relatively early on in the process of idea generation. Similarly, during the screening stage, it is likely that some aspects of the design will become obvious before the phase is finally complete. Therefore, the preliminary work on these parts of the design could commence at that point. This principle can be taken right through all stages, one stage commencing before the previous one has finished, so there is simultaneous or concurrent work on the stages (see Figure 3.10(b)).

We can link this idea with the idea of uncertainty reduction, discussed earlier, when we made the point that uncertainty reduces as the design progresses. This also applies to each stage of innovation. If this is the case, there must be some degree of certainty that the next stage can take as its starting point prior to the end of the previous stage. In other words, designers can continually react to a series of decisions and clues given to them by those working on the preceding stage. However, this can only work if there is effective communication between each pair of stages.

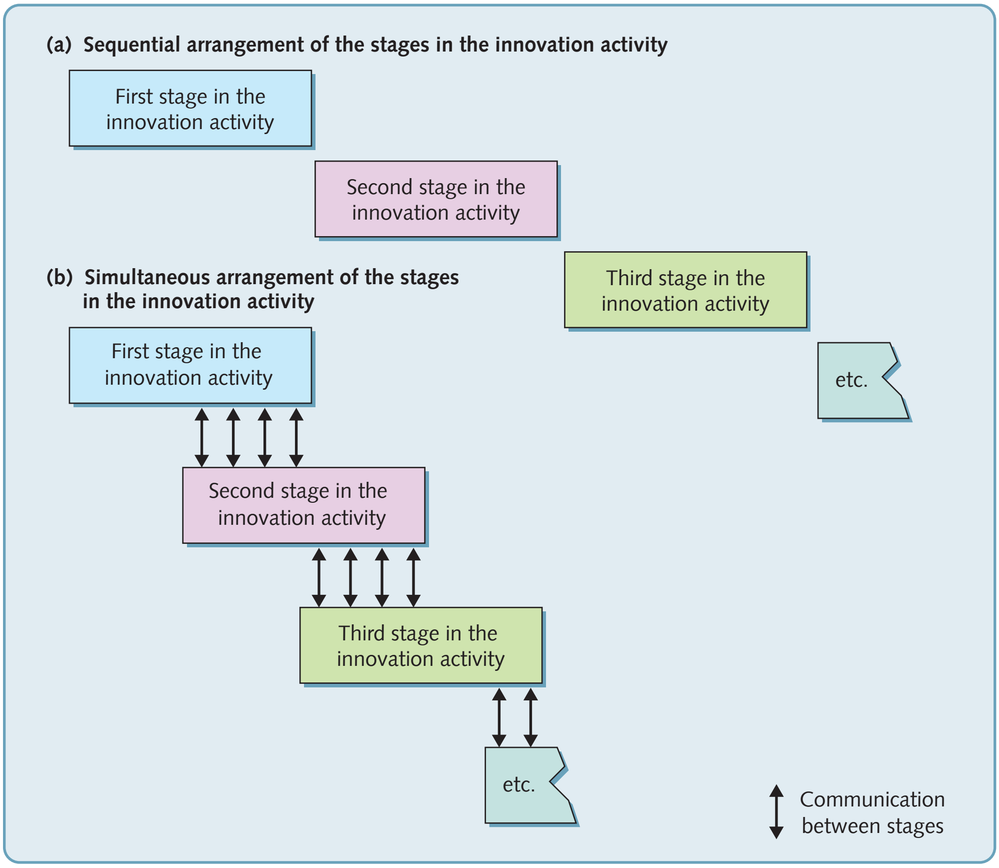

Figure 3.10 **(a) Sequential arrangement of the stages in the innovation activity, (b) Simultaneous arrangement of the stages in the innovation activity**

Deploying strategic intervention and resolving conflicts early

A design decision, once made, need not shape the final offering irrevocably. All decisions can be changed, but it becomes increasingly difficult to do so as the innovation process progresses. At the same time, early decisions are often the most difficult to make because of the high level of uncertainty surrounding what may or may not work as a final design. This is why the level of debate, and even disagreement, over the characteristics of an offering can be at its most heated in the early stages of the process. One approach is to delay decision-making in the hope that an obvious 'answer' will emerge. The problem with this is that if decisions to change are made later in the innovation process, these changes will be more disruptive than if they are made early on. The implication of this is that first, it is worth trying to reach consensus in the early stages of the innovation process even if this seems to be delaying the total process in the short term, and second, there is a particular need for strategic intervention into the innovation process by senior management during these early stages. Unfortunately, there is a tendency for senior managers, after setting the initial objectives of the innovation process, to 'leave the details' to technical experts. They may only become engaged with the process again in the later stages as problems start to emerge that need reconciliation or extra resources.

OPERATIONS PRINCIPLE

The innovation process requires strategic attention early, when there is most potential to affect design decisions.

Organising innovation processes in a way that reflects the nature of the offering

The innovation process will almost certainly involve people from several different areas of the business that will have some part in making the decisions shaping the final offering. Yet any design project will also have an existence of its own. It will have a project name, an individual manager or group of staff who are championing the project, a budget and, hopefully, a clear strategic purpose in the organisation. The organisational question is which of these two ideas – the various organisational functions that contribute to the innovation or the project itself – should dominate the way in which the design activity is managed?

There is a range of possible organisational structures – from pure functional to pure project forms. In a purely functional organisation, all staff associated with the innovation project are based unambiguously in their functional groups. There is no project-based group at all. They may be working full-time on the project but all communication and liaison are carried out through their functional manager. At the other extreme, all the individual members of staff from each function involved in an innovation project could be moved out of their functions and perhaps even co-located in a task force dedicated solely to the project. A project manager who might hold the entire budget allocated to the innovation project could lead the task force. Not all members of the task force necessarily have to stay in the team throughout the design period, but a substantial core might see the project through from start to finish. Some members of a design team may even be from other companies. In-between these two extremes there are various types of matrix organisation with varying emphasis on these two aspects of the organisation (see Figure 3.11):[13]

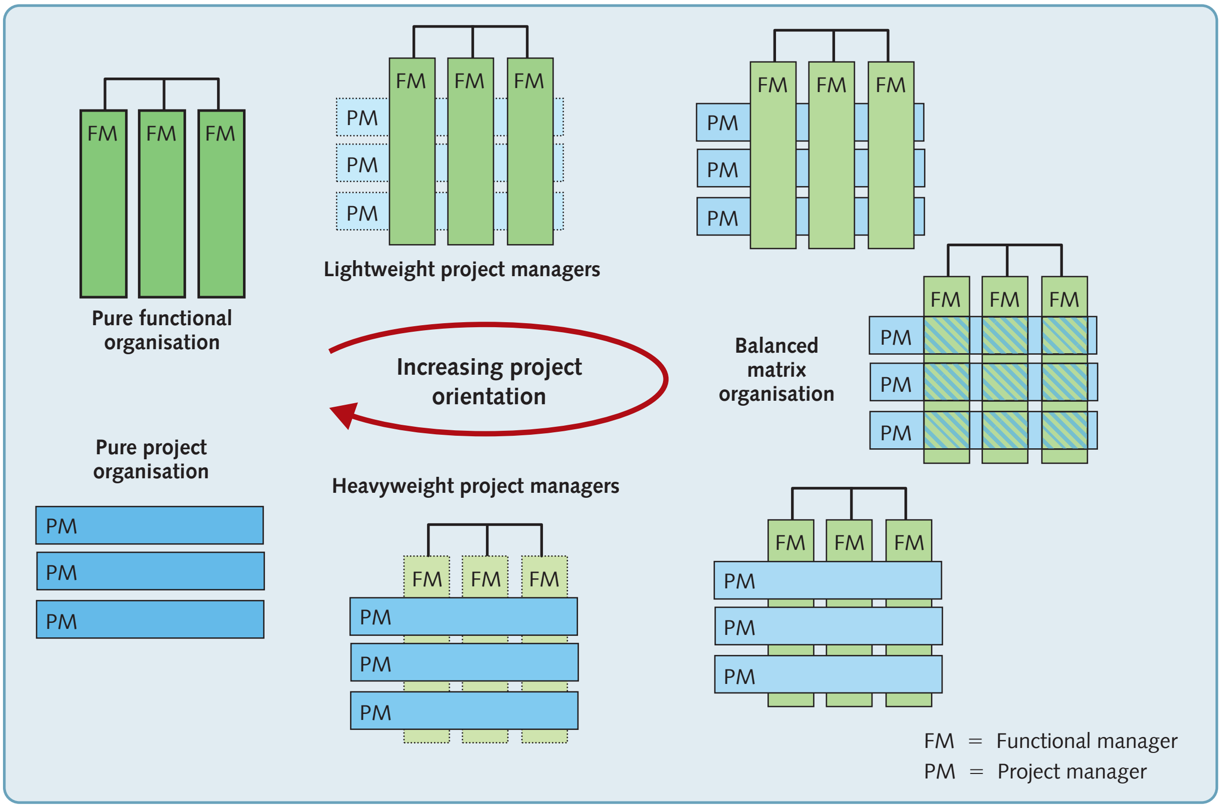

Figure 3.11 Organisation structures for innovation processes

- *Functional organisation*: the innovation project is divided into segments and assigned to relevant functional areas and/or groups within functional areas. The project is co-coordinated by functional and senior management.
- *Functional matrix* (or lightweight project manager): a person is formally designated to oversee the project across different functional areas. This person may have limited authority over the functional staff involved and serves primarily to plan and co-ordinate the project. Functional managers retain primary responsibility for their specific segments of the project.
- *Balanced matrix*: a person is assigned to oversee the project and interacts on an equal basis with functional managers. This person and the functional managers work together to direct innovation activities and approve technical and operational decisions.
- *Project matrix* (or heavyweight project manager): a manager is assigned to oversee the project and is responsible for its completion. Functional managers' involvement is limited to assigning personnel as needed and providing advisory expertise.
- *Project team* (or tiger team): a manager is given responsibility for a project team composed of a core group of personnel from several functional areas assigned on a full-time basis. The functional managers have no formal involvement.

Although there is no clear 'winner' amongst the alternative organisational structures, there is increasing support for structures towards the project rather than functional end of the continuum. However, some authorities argue that heavyweight project manager structures and dedicated project teams are the most efficient forms of organisation in driving competitiveness, shorter lead times and technical efficiency.

Perhaps of more interest is the suitability of the alternative structures for different types of innovation. Matrix structures are generally deemed to be appropriate for both simple and highly complex projects. Dedicated project teams, on the other hand, are seen as appropriate for projects with a high degree of uncertainty, where their flexibility becomes valuable. Functionally based design structures, with resources clustered around a functional specialism, help the development of technical knowledge. Some organisations do manage to capture the deep technological and skills development advantages of functional structures, while at the same time co-coordinating between the functions to ensure satisfactory delivery of new service and product ideas. Perhaps the best known of these organisations is Toyota. It has a strong functionally based organisation to develop its offerings. It adopts highly formalised development procedures to communicate between functions and places strict limits on the use of cross-functional teams. What is really different is their approach to devising an organisational structure for innovation that is appropriate for them. The argument that most companies have adopted to justify cross-functional project teams goes something like this. 'Problems with communication between traditional functions have been the main reasons for failing to deliver new innovation ideas to specification, on time and to budget. Therefore let us break down the walls between the functions and organise resources around the individual development projects. This will ensure good communication and a market-oriented culture'. Toyota and similar companies, on the other hand, have taken a different approach. Their argument goes something like this. 'The problem with cross-functional teams is that they can dissipate the carefully nurtured knowledge that exists within specialist functions. The real problem is how to retain this knowledge on which our future innovation depends, while overcoming some of the traditional functional barriers that have inhibited communication between the functions. The solution is not to destroy the function but to devise the organisational mechanisms to ensure close control and integrative leadership which will make the functional organisation work.'[14]

Critical commentary

The whole process-based approach to innovation could be interpreted as implying that all new offerings are created in response to a clear and articulated customer need. While this is usually the case, especially for services and products that are similar to (but presumably better than) their predecessors, more radical innovations are often brought about by the innovation itself creating demand. Customers don't usually know that they need something radical. For example, in the late 1970s people were not asking for microprocessors; they did not even know what they were. An engineer in the USA improvised them for a Japanese customer who made calculators. Only later did they become the enabling technology for the PC and, after that, the innumerable devices that now dominate our lives. Similarly, fly-by-wire, digital cameras, Maersk's super-slow container ships, sushi on conveyor belts and the iPad are all examples of innovations that have been 'pushed' by firms rather than 'pulled' by pre-existing customer demand.

● Nor does everyone agree with the dominant rational model[15] in which possible design options are progressively reduced stage by stage through the optimisation of known constraints and objectives. For some, this neat model of the innovation, which underlies much business and engineering design literature, fails to accurately reflect the creativity, arguments and chaos that sometimes characterise real innovation projects. First, they argue, managers do not start out with an infinite number of options. No one could process that amount of information – and anyway, designers often have some set solutions in their mind, looking for an opportunity to be used. Second, the number of options being considered often increases as time goes by. This may actually be a good thing, especially if the activity was unimaginatively specified in the first place. Third, the real process of innovation involves cycling back, often many times, as potential solutions raise fresh questions or become dead ends, and as requirements and constraints evolve. In summary, the idea of the design funnel does not describe the process of innovation, nor does it necessarily even describe what should happen. The action-centric or co-evolution[16] perspective of innovation represents the antithesis of the rational model. It posits that offerings are designed through a combination of emotion and creativity; that the process by which this is done is generally improvised; and the sequencing of stages is not universal in innovation processes.

SUMMARY CHECKLIST

- ☐ Is the importance of innovation as a contributor to achieving strategic impact fully understood?
- ☐ Is innovation really treated as a process?
- ☐ Is the innovation process itself designed with the same attention to detail as any other process?
- ☐ Are innovation objectives specified to give a clear priority between quality, speed, dependability, flexibility, cost, and sustainability?
- ☐ Are the stages in the innovation process clearly defined?
- ☐ Are ideas and concepts for new offerings captured from all appropriate internal and external sources?
- ☐ Are potential offerings screened in a systematic manner in terms of their feasibility, acceptability and vulnerability?
- ☐ During preliminary design, have all possibilities for design standardisation, commonality, and modularisation of design elements been explored?
- ☐ Has the concept of mass customisation been linked to the innovation process?
- ☐ Are potential offerings thoroughly evaluated and tested during the innovation process?
- ☐ Are the resources for developing innovation adequate?
- ☐ Is sufficient capacity devoted to the innovation process?
- ☐ Have all options for outsourcing parts of the innovation process been explored?
- ☐ Has the possibility of involving customers formally in the innovation process been explored?
- ☐ Are appropriate technologies such as CAD, digital twinning, and knowledge management, being used in the innovation process?
- ☐ Are the design of offering and the design of processes that create and deliver it considered together as one integrated process?
- ☐ Is overlapping (simultaneous) of the stages in the innovation process used?
- ☐ Is senior management effort deployed early enough to ensure early resolution of design conflict?
- ☐ Does the organisational structure of the innovation process reflect the nature of the offering?
- ☐ Are some functions of the business more committed to innovating new service and product offerings than others?
- ☐ If so, have the barriers to cross-functional commitment been identified and addressed?

CASE STUDY

Developing 'Savory Rosti-crisps' at Dreddo Dan's

'*Most people see the snack market as dynamic and innovative, but actually it is surprisingly conservative. Most of what passes for innovation is in fact tinkering with our marketing approach, things like special offers, promotion tie-ins, and so on. We occasionally put new packs round our existing products and even more occasionally we introduce new flavors in existing ranges. Rarely though does anyone in this industry introduce something radically different. That is why "Project Orlando" is both exciting and scary*.'

Monica Allen, the technical vice-president of PJT's snack division, was commenting on a new product that will be marketed under PJT's best-known brand 'Dreddo Dan's Surfer Snacks'. The Dreddo Dan's brand made use of surfing and outdoor 'action oriented youth' imagery but, in fact, was aimed at a slightly older generation, who, although aspiring to such a lifestyle, had more discretionary spend for the premium snacks in which the brand specialised. Current products marketed under the brand included both fried and baked snacks in a range of exotic flavours. The project, internally known as Project Orlando, was a baked product that had been 'in development' for almost 3 years but had hitherto been seen very much as a long-term development, with no guarantee of it ever making it through to market launch. PJT had several of these long-term projects running at any time. They were allocated a development budget, but usually no dedicated resources were associated with the project. Less than half of these long-term projects ever even reached the stage of being test marketed. Around 20 per cent did not get past the concept stage, and less than 20 per cent ever went into production. However, the company viewed the development effort put into these 'failed' products as being worthwhile because it often led to 'spin-off' developments and ideas that could be used elsewhere. Up to this point, 'Orlando' had been seen as unlikely ever to reach the test marketing stage, but that had now changed dramatically.

'Orlando' was a concept for a range of snack foods, described within the company as 'savory potato cookies'. Essentially, they were one and a half inch discs of crisp, fried potato with a soft dairy cheese-like filling. The idea of incorporating dairy fillings in snacks had been discussed within the industry for some time, but the problems of manufacturing such a product were formidable. Keeping the product crisp on the outside yet soft in the middle, while at the same time ensuring microbiological safety, would not be easy. Moreover, such a product would have to be capable of being stored at ambient temperatures, maintain its physical robustness and have a shelf life of at least 3 months.

Bringing Orlando products to market involved overcoming three types of technical problem. First, the formulation and ingredient mix for the product had to maintain the required texture yet be capable of being baked on the company's existing baking lines. The risk of developing an entirely new production technology for the offering was considered too great. Second, extruding the mixture into baking moulds, while maintaining microbiological integrity (dairy products are difficult to handle) would require new extrusion technology. Third, the product would need to be packaged in a material that both reflected its brand image but also kept the product fresh through its shelf life. Existing packaging materials were unlikely to provide sufficient shelf life. The first of these problems had, more or less, been solved in PJT's development laboratories. The second two problems now seemed less formidable because of a number of recent technological breakthroughs made by equipment suppliers and packaging manufacturers. This had convinced the company that Orlando was worth significant investment and the company's board had given it priority development status. Even so, it was not expected to come to the market for another 2 years and was seen by some as potentially the most important new product development in the company's history.

The project team

Immediately after the board's decision, Monica had accepted responsibility to move the development forward. She decided to put together a dedicated project team to oversee the development. '*It is important to have representatives from all relevant parts of the company. Although the team will carry out much of the work themselves, they will still need the cooperation and the resources of their own departments. So, as well as being part of the team,* *they are also gateways to expertise around the company.'* The team consisted of representatives from marketing, the development kitchens (laboratories), PGT's technology centre (a development facility that served the whole group, not just the snack division), packaging engineers and representative from the division's two manufacturing plants. All but the manufacturing representatives were allocated to the project team on a full-time basis. Unfortunately, manufacturing had no one who had sufficient process knowledge and who could be spared from their day-to-day activities.

Development objectives

Monica had tried to set the objectives for the project in her opening remarks to the project team when they had first come together. '*We have a real chance here to develop an offering that not only will have major market impact, but will also give us a sustainable competitive advantage. We need to make this project work in such a way that competitors will find it difficult to copy what we do. The formulation is a real success for our development people, and as long as we figure out how to use the new extrusion method and packaging material, we should be difficult to beat. The success of Orlando in the marketplace will depend on our ability to operationalise and integrate the various technical solutions that we now have access to. The main problem with this type of offering is that it will be expensive to develop and yet, once our competitors realise what we are doing, they will come in fast to try and out-innovate us. Whatever else we do, we must ensure that there is sufficient flexibility in the project to allow us to respond quickly when competitors follow us into the market with their own 'me-too' products. We are not racing against the clock to get this to market, but once we do make a decision to launch, we will have to move fast and hit the launch date reliably. Perhaps most important, we must ensure that the crisps are 200 per cent safe. We have no experience in dealing with the microbiological testing that dairy-based food manufacture requires. Other divisions of PJT do have this experience and I guess we will be relying heavily on them.*' (Monica Allen)

Monica, who had been tasked with managing the, now much expanded, development process had already drawn up a list of key decisions she would have to take:

- How to resource the innovation project – The division had a small development staff, some of whom had been working on Project Orlando, but a project of this size would require extra staff, amounting to about twice the current number of people dedicated to the innovation process.
- Whether to invest in a pilot plant – The process technology required for the new project would be unlike any of the division's current technology. Some companies in the frozen food industry used similar technology and one option would be to carry out trials at these (non-competitor) companies' sites. Alternatively, the Orlando team could build their own pilot plant, which would enable them to experiment in-house. As well as the significant expense involve, this would raise the problem of whether any process innovations would work when scaled-up to full size. However, it would be far more convenient for the project team and allow them to 'make their mistakes' in private.
- How much development to outsource – Because of the size of the project, Monica had considered outsourcing some of the innovation activities. Other divisions within the company may be able to undertake some of the development work and specialist consultancies operated in the food processing industries. The division had never used any of these consultancies before but other divisions had occasionally done so.
- How to organise the innovation activities – Currently the small development function had been organised around loose functional specialisms. Monica wondered whether this project warranted the creation of a separate department independent of the current structure. This might signal the importance of this innovation project to the whole division.

Fixing the budget

The budget to develop Project Orlando through to launch had been set at $30 million. This made provision to increase the size of the existing innovation team by 70 per cent over a 20-month period (for launch 2 years from now). It also included enough funding to build a pilot plant, which would allow the team the flexibility to develop responses to potential competitor reaction after the launch. Of the $30m, around $18m was for extra staff and contracted-out innovation work, $7.5m for the pilot plant and $4.5m for one-off costs (such as the purchase of test equipment, etc.). Monica was unsure whether the budget would be big enough.

'*I know everyone in my position wants more money, but it is important not to under fund a project like this. Increasing our development staff by 70 per cent is not really enough. In my opinion, we need an increase of at least 90 per cent to make sure that we can launch when we want. This would need another $5m, spread over the next 20 months. We could get this by not building the pilot plant*

I suppose, but I am reluctant to give that up. It would mean begging for test capacity on other companies plants, *which is never satisfactory from a knowledge-building viewpoint. Also it would compromise security. Knowledge of what we were doing could easily leak to competitors. Alternatively, we could subcontract more of the research, which may be less expensive, especially in the long run; but I doubt if it would save the full $5m we need. More important, I am not sure that we should subcontract anything that would compromise safety, and increasing the amount of work we send out may do that. No, it's got to be the extra cash or the project could overrun. The profit projections for the Orlando products look great (see Table 3.2), but delay or our inability to respond to competitor pressures would depress those figures significantly. Our competitors could get into the market only a little after us. Word has is that Marketing's calculations indicate a delay of only six months could not only delay the profit stream by six months but also cut it by up to 30%.'*

Monica was keen to explain two issues to the management committee when it met to consider her request for extra funding. First, that there was a coherent and well-thought-out strategy for the innovation project over the next two years. Second, that saving $5m on Project Orlando's budget would be a false economy.

Table 3.2 Preliminary 'profit stream' projections for the Project Orlando offering, assuming launch in 24 months' time

Time period*	1	2	3	4	5	6	7
Profit flow ($ million)	10	20	50	90	120	130	135

*6 month periods

QUESTIONS

1 How would you rank the innovation objectives for the project?

2 What are the key issues in resourcing this innovation process?

3 What are the main factors influencing the resourcing decisions?

4 What advice would you give Monica?

APPLYING THE PRINCIPLES

Some of these exercises can be answered by reading the chapter. Others will require some general knowledge of business activity and some might require an element of investigation. **Model answers for the first two questions are available on the student companion website.** For model answers to the other questions in this section, please ask your tutor/lecturer.

1 One product for which customers value a very wide range of product types is that of domestic paint. Most people like to express their creativity in their choice of paints and other home decorating products that they use in their homes. Clearly, offering a wide range of paint must have serious cost implications for the companies that manufacture, distribute and sell the product. Visit a store that sells paint and get an idea of the range of products available on the market. How do you think paint manufacturers and retailers could innovate to increase variety but minimise costs?

2 'We have to get this new product out fast', said the Operations Director. 'Our competitors are close behind us and I believe their products will be almost as good as ours when they launch them'. She was talking about a new product that the company hoped would establish them as the leader in the market. The company had put together a special development team, together with their own development laboratory. They had spent £10,000 on equipping the laboratory and the cost of the development engineers would be £20,000 per quarter. It was expected that the new product would be fully developed and ready for launch within 6 quarters. It would be done through a specialist agency that charged £10,000 per quarter and would need to be in place two quarters prior to the launch. If the company met their launch date, it was expected that they could charge a premium price that would result in profits of approximately £50,000 per quarter. Any delay in the launch would result in a reduction in profits to £40,000 per quarter. If this development project were delayed by two quarters, how far would the break-even point for the project be pushed back?

3 Innovation becomes particularly important at the interface between offerings and the people that use them. Consider two types of website,

(a) Those which are trying to sell something such as Amazon.com, and

(b) Those primarily concerned with giving information, for example reuters.com or nytimes.com.

What constitutes good innovation for these two types of website? Find examples of particularly good and particularly poor web design and explain the issues you've considered in making the distinction between them.

4 According to the Ellen MacArthur Foundation, a circular economy is '*one that is restorative and regenerative by design, and which aims to keep products, components and materials at their highest utility and value at all times, distinguishing between technical and biological cycles*'. See also the example earlier in this chapter. What do you see as the main barriers to a more widespread adoption of the idea?

5 A janitor called Murray Spangler invented the vacuum cleaner in 1907. One year later he sold his patented idea to William Hoover, whose company went on to dominate the market. Now, the Dyson vacuum cleaner has jumped from nothing to a position where it dominates the market. The Dyson product dates back to 1978 when James Dyson designed a cyclone-based cleaner. It took 5 years and 5,000 prototypes before he had a working design. However, existing vacuum cleaner manufacturers were not as impressed – two rejected the design outright. So Dyson started making his new design himself and within a few years Dyson cleaners were outselling the rivals who had once rejected them. The aesthetics and functionality of the design help to keep sales growing in spite of a higher retail price. To Dyson good design '*is about looking at everyday things with new eyes and working out how they can be made better. It's about challenging existing technology*'.

(a) What was Spangler's mistake?

(b) What do you think makes 'good design' in markets such as the domestic appliance market?

(c) Why do you think the two major vacuum cleaner manufacturers rejected Dyson's ideas?

(d) How did design make Dyson a success?

6 It sounds like a joke, but it is a genuine product innovation. It's green, it's square and it comes originally from Japan. It's a square watermelon. Why square? Because Japanese grocery stores are not large and space cannot be wasted. Similarly, a round watermelon does not fit into a refrigerator very conveniently. There is also the problem of trying to cut the fruit when it keeps rolling around. An innovative Japanese farmer solved the problem with the idea of making a cube-shaped watermelon that could easily be packed and stored. There is no genetic modification or clever science involved in growing watermelons. It simply involves placing the young fruit into wooden boxes with clear sides. During its growth, the fruit naturally swells to fill the surrounding shape. (a) Why is a square watermelon an advantage? (b) What does this example tell us about product design?

Notes on chapter

1 Trott, P. (2011) *Innovation Management and New Product Development*, Pearson Education

2 Sources include: Heisler, Y. (2013) 'Apple's iPhone: The untold story', *Network World,* 1 September; Vogelstein, F. (2013) 'And Then Steve Said, "Let There Be an iPhone", *New York Times Magazine,* 4 October; Ahmed, M. (2012) 'Apple nearly ditched the iPhone, designer admits', *The Times,* 31 July; Breillatt, A. (2012) 'You Can't Innovate Like Apple', www.pragmaticmarketing.com

3 Sources include: *The Economist* (2012) 'The last Kodak moment?', 14 January; Wright, M. (2010) 'Kodak develops: A film giant's self-reinvention', *Wired Magazine,* 15 February.

4 Henderson, R. M. and Clark, K. B. (1990) 'Architectural innovation: The reconfiguration of existing product technologies and the failure of established firms', *Administrative Science Quarterly,* vol. 35, no. 1, 9–22.

5 Goodwin, L. (2015) 'How to bust the biggest myths about the circular economy', *The Guardian,* 12 March; Clegg, A. (2015) 'Sustainable innovation: shaped for the circular economy', *Financial Times,* 26 August; Company website, Newlife Paints, http://www.newlifepaints.com/

6 The term was originally used by von Hippel, E. (1986) 'Lead users: a source of novel product concepts,' *Management Science,* vol. 32, no. 7 (July)

7 Definition from Financial Times Lexicon, lexicon.ft.com/

8 Anderson, E., Lin, S., Simester, D. and Tucker, C. (2015) 'Harbingers of failure', *Journal of Marketing Research*: October 2015, vol. 52, no. 5,

9 Mass customisation was first fully articulated in Pine, B.J. (1993) *Mass Customization: the New Frontier in Business Competition,* Harvard Business School Press.

10 Coppinger, R. (2016)' Digital twins: CAD design through the looking glass', *Engineering and Technology Magazine,* November 9

11 Volkmann, D. (2016) 'The Rise of Digital Twins', Linkedin blog, linkedin.com/pulse, 7 November.

12 Wheelwright, S.C. and Clark, K.B. (1995) *Leading Product Development*, The Free Press.

13 This idea is based on one presented by Hayes, Wheelwright and Clark, in Hayes, R.H., Wheelwright, S.C. and Clark , K.B. (1988) *Dynamic Manufacturing,* The Free Press.

14 Sobek, D.K. II, Liker, J.K. and Ward, A.K. (1998) 'Another look at how Toyota integrates product development', *Harvard Business Review,* July-August.

15 This term is coined by Pahl, G., and Beitz, W. in (1996) *Engineering Design: A systematic Approach,* Springer-Verlag.

16 From Dorst, K. and Cross, N. (2001) Creativity in the design process: Co-evolution of problem-solution,' *Design Studies* (22), September 2001, 425–437.

TAKING IT FURTHER

Bangle, C. (2001) 'The Ultimate Creativity Machine: How BMW turns art into profit', *Harvard Business Review, Jan, 47–55.* *A good description of how good aesthetic design translates into business success.*

Bruce, M. and Bessant, J. (2002) *Design In Business: Strategic innovation through design, Financial Times Prentice Hall and The Design Council.* *Probably one of the best overviews of innovation in a business context available today.*

Chesbrough, H.W. (2003) *Open Innovation: The New Imperative for Creating and Profiting from Technology,* Harvard Business School Press. *A good overview of the reasons for open innovation and the challenges of making it work.*

Christensen, C M. (1997) *The Innovator's Dilemma*, *Harvard Business School Press.* *Ground-breaking booking on disruptive innovation*

Kahney, L. (2014) *Jony Ive: The Genius Behind Apple's Greatest Products, Penguin*. *Inside the mind of one of the World's great designers.*

Mckeown, M. (2008) *The Truth About Innovation*, *Prentice Hall.* *A practical guide to innovation, including lots of examples of success and failure.*

Nambison, S. and Sawhney, M. (2007) 'A Buyer's Guide to the Innovation Bazaar', *Harvard Business Review, June.* *Provocative innovation ideas.*

Reason, B., Løvlie, L. and Brand Flu, M. (2016) *Service Design for Business: A Practical Guide to Optimizing the Customer Experience, John Wiley & Sons.* *A book that looks at the 'entire service experience' and very readable.*

4 Operations scope and structure

Introduction

Both the structure and the scope of an operation's supply network are decisions that shape how the operation interacts with its markets, its suppliers, in fact with the world in general. No operation exists in isolation. All operations are part of a larger and interconnected network of operations. This is called the operation's *supply network*. It will include the operation's suppliers and customers, as well as suppliers' suppliers and customers' customers, and so on. This chapter is concerned with the key strategic decisions that influence the scope and structure of this network, and more specifically the role of the individual operations within it. At a strategic level, operations managers are involved in deciding how much of the network it should own. This is called the *scope* of the operation. Put another way, the scope of the operation defines what it is going to do itself and what it will buy-in from suppliers. They are also concerned with the shape and form of their network. This is called the *structure* of the network. It involves deciding (or at least, influencing) its overall shape, the location of each operation and how big the parts of the network that it owns should be. This chapter treats the issues related to both the scope and structure decisions, and because both scope and structure decisions require an estimate of future demand, this chapter includes a supplement on forecasting. (*See* Figure 4.1).

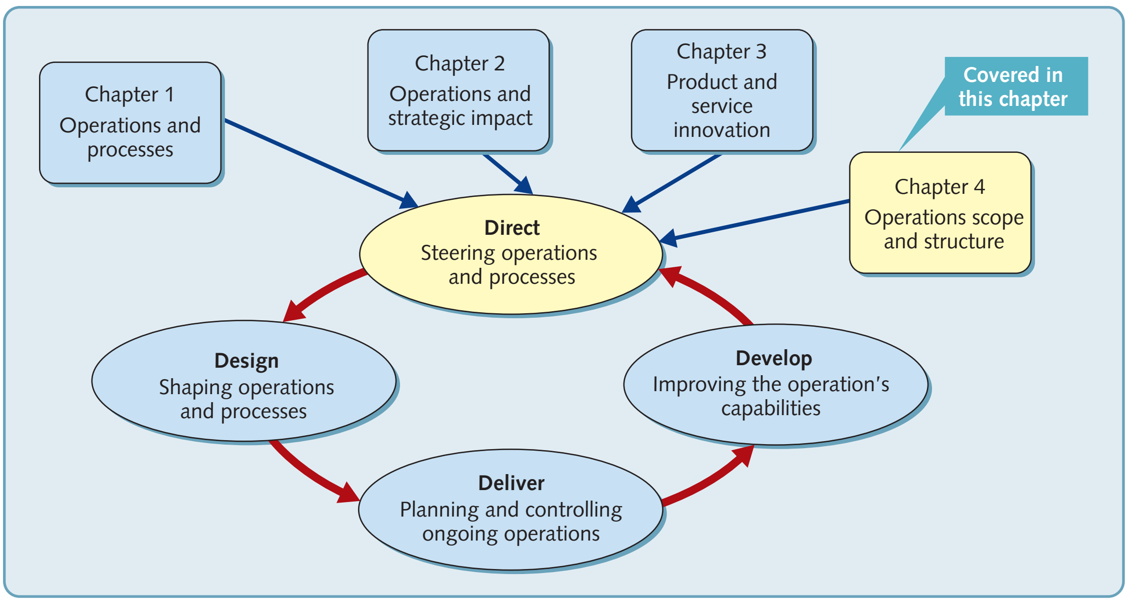

Figure 4.1 This chapter covers the structure and scope of operations

EXECUTIVE SUMMARY

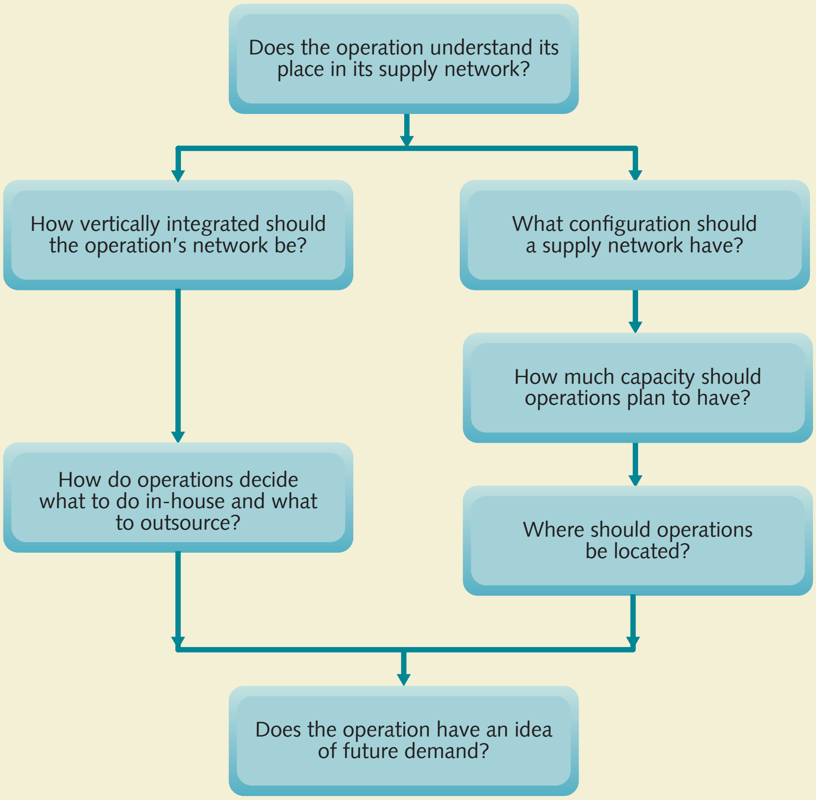

Does the operation understand its place in its supply network?

The supply network includes the chains of suppliers providing inputs to the operation, the chain of customers who receive outputs from the operation and sometimes other operations that may at times compete and other times cooperate. To understand its place and role, an operation must have an idea of how it fits into the scope and structure of its supply network. The 'structure' of an operation's supply network relates to the shape and form of the network. The scope of an operation's supply network relates to the extent that an operation decides to do the activities performed by the network itself, as opposed to requesting a supplier to do them. The structure and scope of an operation's supply network is important because it helps an understanding of competitiveness, it helps identify significant links in the network, and it helps focus on long-term issues.

How vertically integrated should the operation's network be?

The scope to which an operation controls its supply network is the extent that it does things itself, as opposed to relying on other operations to do things for it. This is often referred to as 'vertical integration'. An organisation's vertical integration strategy can be defined in terms of the direction of integration, the extent of integration and the balance among the vertically integrated stages. The decision as to whether to vertically integrate is largely a matter of a business balancing the advantages and disadvantages as they apply to them.

How do operations decide what to do in-house and what to outsource?

Outsourcing is, 'an arrangement in which one company provides services for another company that could also be, or usually have been, provided in-house'. The difference between vertical integration and outsourcing is largely a matter of scale and direction. Like the vertical integration decision, it is often a matter of balancing advantages against disadvantages under particular circumstances. Assessing the advisability of outsourcing should also include consideration of the strategic importance of the activity and the operation's relative performance.

What configuration should a supply network have?

Even when an operation does not directly own other operations in its network, it may still wish to change the shape of the network by reconfiguring it, in order to change the nature of the relationships. A number of trends are reshaping networks in many industries. These include reducing the number of individual suppliers, the disintermediation of some parts of the network and a greater tolerance of other operations being both competitors and complementors at different times (co-opertition). An idea that is closely related to that of co-opertition in supply networks is that of the 'business ecosystem', defined as: An economic community supported by a foundation of interacting organisations and individuals.

How much capacity should operations plan to have?

The amount of capacity an organisation will have depends on its view of current and future demand. It is when its view of future demand is different from current demand that this issue becomes important. When an organisation has to cope with changing demand, a number of capacity decisions need to be taken. These include choosing the optimum capacity for each site, balancing the various capacity levels of the operation in the network, and timing the changes in the capacity of each part of the network. Important influences on these decisions include the concepts of economy and diseconomy of scale.

Where should operations be located?

When operations change their location, their assumption is that the potential benefits of a new location will outweigh any cost and disruption involved in changing location. When operations do move, it is usually because of changes in demand and/or changes in supply. The factors that determine a location are such things as labour, land and utility costs, the image of the location, its convenience for customers and the suitability of the site itself.

DIAGNOSTIC QUESTION

Does the operation understand its place in its supply network?

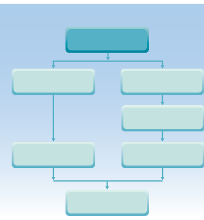

The scope of an operation's supply network relates to the extent that an operation decides to do the activities performed by the network itself, as opposed to requesting a supplier to do them. The 'structure' of an operation's supply network relates to the shape and form of the network. But before we examine these issues, we need to establish what we mean by 'a supply network' and why it is important that operations managers understand their position in it.

'A supply network is an interconnection of organizations that relate to each other through upstream and downstream linkages between the different processes and activities that produce value in the form of products and services to the ultimate consumer'. [1] In other words, a supply network is the means of setting an operation in the context of all the other operations with which it interacts, some of which are its suppliers and its customers. Materials, parts, other information, ideas and sometimes people all flow through the network of customer–supplier relationships formed by all these operations.

Terminology is important when describing supply networks. On its supply side an operation has its suppliers of parts, or information, or services. These suppliers themselves have their own suppliers who in turn could also have suppliers, and so on. On the demand side the operation has customers. These customers might not be the final consumers of the operation's products or services; they might have their own set of customers. Suppliers that directly supply the operation are often called first-tier suppliers. Second-tier suppliers supply them; however, some second-tier suppliers, as well as supplying first-tier suppliers, may also supply an operation directly, thus missing out a link in the network. Similarly, 'first-tier' customers are the main customer group for the operation. These in turn supply 'second-tier' customers, although again the operation may at times supply second-tier customers directly. The suppliers and customers who have direct contact with an operation are called its immediate supply network. Figure 4.2 illustrates this.

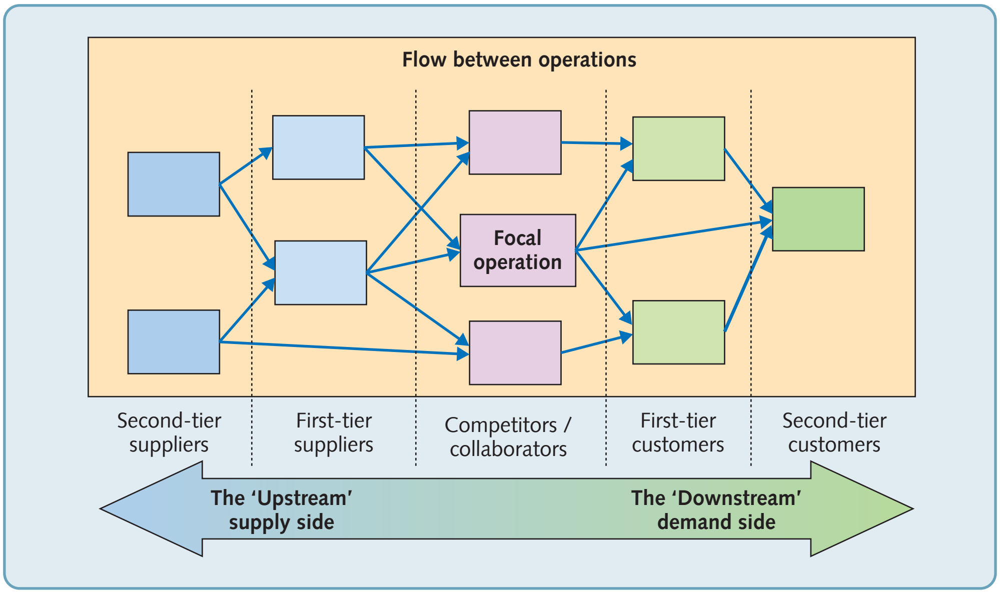

Figure 4.2 **Supply network terminology**

Two-way flow through the network

Materials, parts, information, ideas and sometimes people all flow through the network of customer–supplier relationships formed by all these operations. But also along with the forward flow of transformed resources (materials, information and customers) in the network, each customer–supplier linkage will feed back orders and information. For example, when stocks run low, retailers place orders with distributors, who likewise place orders with the manufacturer; in turn the manufacturer will place orders with its suppliers, who will replenish their own stocks from their own suppliers. So flow is a two-way process with items flowing one way and information flowing the other.

Yet, it is not only manufacturers who are part of a supply network. The flow of physical materials may be easier to visualise, but service operations also have suppliers and customers who themselves have their own suppliers and customers. One way to visualise the supply networks of some service operations is to consider the downstream flow of information that passes between operations. Most financial service supply networks can be thought about like this. However, not all service supply networks deal primarily in information. For example, property companies that own and/or run shopping malls have suppliers who provide security services, cleaning services, maintenance services, and so on. These first-tier suppliers will themselves receive service from recruitment agencies, consultants, and so on. First-tier customers of the shopping mall are the retailers who lease retail space within the mall, who themselves serve retail customers. This is a supply network like any other. What is being exchanged between operations is the quality, speed, dependability, flexibility and cost of the services each operation supplies to its customers. In other words, there is a flow of 'operations performance' through the network. And although visualising the flow of 'performance' through supply networks is an abstract approach to visualising supply networks, it is a unifying concept. Broadly speaking, all types of supply network exist to facilitate the flow of 'operations performance'.

Here are two examples of where supply network issues have been particularly important to the businesses (and people) involved.

EXAMPLE

Contrasting strategies on structure and scope: ARM versus Intel[2]

Nothing better illustrates the idea that there is more than one approach to competing in the same market than the contrasting business models of ARM and Intel in the microchip business. At one point, ARM's chip designs were to be found in almost 99 per cent of mobile devices in the world, while Intel dominates the PC and server markets. Yet ARM and Intel are very different companies, with different approaches to the structure and scope of their operations and, some claim, very different prospects for their future. They certainly are of a different size. In revenue terms, Intel was around 50 times bigger than ARM. More interestingly, Intel is vertically integrated, both designing and manufacturing its own chips, while ARM is essentially a chip designer, developing intellectual property. It then licenses its processor designs to manufacturers such as Samsung, who in turn rely on subcontracting 'chip foundry' companies to do the actual manufacturing (ironically, including Intel).

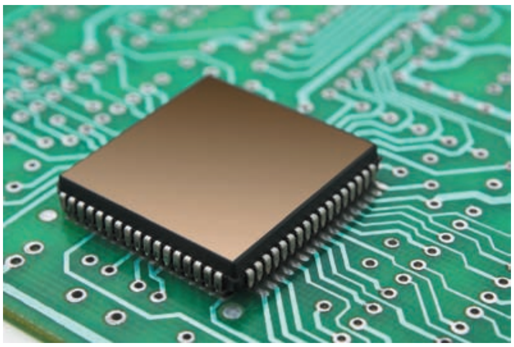

Intel's integrated supply network monitors and controls all stages of production, from the original design concept right through to manufacturing. Keeping on top of fast-changing and hugely expensive (it can cost around $5 billion to build a new chip-making plant) operations requires very large investments. It is Intel's near monopoly (therefore high volume) of the server and PC markets that helps it to keep its unit prices high, which in turn gives it the ability to finance the construction of the latest semiconductor manufacturing equipment before its competitors. And having the latest manufacturing technology is important; it can mean faster, smaller and cheaper chips with lower power consumption. As one industry source put it, 'Intel is one of the few companies left with the financial resources to invest in state-of-the-art manufacturing research and development. Everyone else – including all the ARM licensees – have to make do with shared manufacturing, mainstream technology, and less-aggressive physics.' By contrast, ARM's supply network strategy was a direct result of their early lack of cash. It did not have the money to invest in its own manufacturing facilities (or to take the risk of subcontracting manufacturing), so it focused on licensing its 'reference designs'. Reference designs provide the 'technical blueprint' of a microprocessor that third parties can enhance or modify as required. This means that partners can take ARM reference designs and integrate them in flexibly to produce different final designs. And over the years a whole 'ecosystem' of tools has emerged to help developers build applications around the ARM design architecture. The importance of ARM's supply 'ecosystem' should not be underestimated. It is an approach that allows ARM's partners to be part of ARM's success rather than cutting them out of revenue opportunities.

EXAMPLE

Disaster at Rana Plaza[3]

It was a disaster that grabbed the attention of the world; on 24 April 2013 the Rana Plaza clothing factory near Dhaka in Bangladesh collapsed, killing at least 1,100 people. At that time many well-known clothing brands were sourcing products, either directly or indirectly, from the factory. It was claimed that local police and an industry association issued a warning that the building was unsafe, but the owners had responded by threatening to fire people who refused to carry on working as usual. Understandably, there was an immediate call for tighter regulation and oversight by the Bangladeshi authorities, and for the predominantly western retailers who sourced from the Rana Plaza, and similar unsafe factories, to accept some of the responsibility for the disaster and change their buying policies. Campaigning organisations including 'Labour Behind the Label', 'War on Want' and 'Made in Europe' urged retailers to be more transparent about their supply chains. They also called for compensation to be paid. But a year after the tragedy the compensation initiative that intended to raise $40 million had raised only $15 million, despite being backed by the UN's International Labour Organisation. Less than half the brands linked to clothes-making at the building had made donations. Benetton and Matalan said they preferred to support other funds that assisted victims, while the French retailer Auchan claimed that they had no official production taking place in the building when it collapsed, so they did not need to contribute towards compensation. Other contributions were relatively small. Walmart, the largest retailer in the world, offered to contribute about $1 million compared to more than $8 million from the far smaller Primark. The Bangladeshi authorities also came in for international criticism. For years they had made only relatively weak attempts to enforce national building regulations, especially if the landlords involved were politically well connected. After the disaster, they promised to apply the laws more rigorously, but such promises had been made before.

What are the options for western retailers? One option is to carry on as before and simply source garments from wherever is cheapest. Doing so would obviously be ethically questionable, but would it also carry a reputational cost? Or would consumers not enquire too deeply about where garments came from if they were cheap enough? Alternatively, retailers could stop sourcing from Bangladesh until there were improvements in the country's garments manufacturing industry. But that may be difficult to enforce unless they took on the responsibility to police the whole supply chain right back to the cotton growers. It would also damage all Bangladeshi firms, even those who try to abide by safety rules. This in turn could also be damaging to the retailers' reputations. The third option is to stay and try to change how things are done in the country. Even before the Rana Plaza disaster, retailers had met with some interested parties and governments to develop a strategy to improve safety in Bangladesh's 5,000 factories. Some individual retailers had launched initiatives. Walmart had launched a fire-safety training academy and Gap had announced a plan to help factory owners to upgrade their plants. However, individual initiatives are no substitute for properly coordinated safety improvements. And anyway, some claim, what right have western companies to impose their rules on another sovereign state?

What do these two examples have in common?

At first glance these two examples seem to be about very different things – although all the issues are covered in this chapter. Yet, the issues exposed by the contrasting vertical integration strategies of ARM and Intel, and the Rana Plaza disaster are connected. The ARM strategy of limiting its activities to a relatively small part of its total supply network (compared with Intel) has proved very successful for the company. Nor is it alone. Many firms have shrunk their 'span of ownership' within their supply network. This is why the idea of 'outsourcing' non-core activities has been so well accepted over the last few decades. Focusing on a relatively narrow range of activities like ARM has done clearly has advantages. Yet outsourcing activities also carries risks; often linked to a loss of control over suppliers. The Rana Plaza example provides an appalling example of how a loss of control over suppliers can lead to both human and reputational tragedy.

Why is it important to consider the whole supply network?

What is undeniable in these two examples is that supply network issues, and the operations' positions in them, had a significant effect on their strategic performance. In addition, there are other reasons why it is important to stand back and look at the whole of (or a large part of) a supply network rather than an individual operation.

It helps an understanding of competitiveness

Immediate customers and immediate suppliers, quite understandably, are the main concern for companies. Yet sometimes they need to look beyond these immediate contacts to understand *why* customers and suppliers act as they do. Any operation has only two options if it wants to understand its ultimate customers' needs. It can rely on all the intermediate customers and customers' customers, and so on, that separate it from its end customers, or it can take responsibility for looking beyond its immediate customers and suppliers itself. Simply relying only on one's immediate network is arguably putting too much faith in someone else's judgement of things which are central to an organisation's own competitive health.

It helps identify significant links in the network

Not everyone in a supply network has the same degree of influence over the performance of the network as a whole. Some operations contribute more to the performance objectives that are valued by end customers. So an analysis of networks needs an understanding of the downstream and the upstream operations that contribute most to end-customer service. For example, the important end customers for domestic plumbing parts and appliances are the installers and service companies who deal directly with consumers. They are supplied by 'stock holders' who must have all parts in stock and deliver them fast. Suppliers of parts to the stock holders can best contribute to their end customers' competitiveness partly by offering a short delivery lead time but mainly through dependable delivery. The key players in this example are the stock holders. The best way of winning end-customer business in this case is to give the stock holder prompt delivery, which helps keep costs down while providing high availability of parts.

OPERATIONS PRINCIPLE

A supply network perspective helps to make sense of competitive, relationship, and longer-tem operations issues.

It helps focus on long-term issues

There are times when circumstances render parts of a supply network weaker than its adjacent links. High street music stores, for example, have been largely displaced by music streaming and downloading services. A long-term supply network view would involve constantly examining technology and market changes to see how each operation in the supply network might be affected.

Scope and structure

The scope and structure of an operation's supply network are strongly related (which is why we treat them together). For example, suppose that a company that runs a shopping mall is dissatisfied with the service it is receiving from its supplier of security services. Also suppose that it is considering three alternatives. Option 1 is to switch suppliers and award the security contract to a different security services supplier. Option 2 is to accept an offer from the company that supplies cleaning services to supply both security and cleaning services. Option 3 is to take over responsibility for security itself; hiring its own security staff who would be put on the mall's payroll. These options are illustrated in Figure 4.3. The first of these options changes neither the structure, nor the scope of this part of the supply network. The shopping mall still has three suppliers and is doing exactly what it did before. All that has changed is that security services are now being provided by another (hopefully better) supplier. However, option 2 changes the structure of the supply network (the mall now has only two suppliers, the combined cleaning and security supplier, and the maintenance supplier), but not the scope of what the mall does (it does exactly what it did before). Option 3 changes both the structure of the network (again, the mall has only two suppliers; cleaning and maintenance services) and the scope of what the mall does (it now also takes on responsibility for security itself).

A further point to make is that both scope and structure decisions actually comprise a number of other 'constituent' decisions.

The scope of an operation's activities within the network is determined by two decisions:

1. The extent and nature of the operation's vertical integration.
2. The nature and degree of outsourcing it engages in.

The structure of an operation's supply network is determined by three sets of decisions:

1. How should the network be configured?
2. What physical capacity should each part of the network have – the long-term capacity decision?
3. Where should each part of the network be located – the location decision?

Note, however, that all of these decisions rely on forecasts of future demand that the supplement to this chapter explores in more detail.

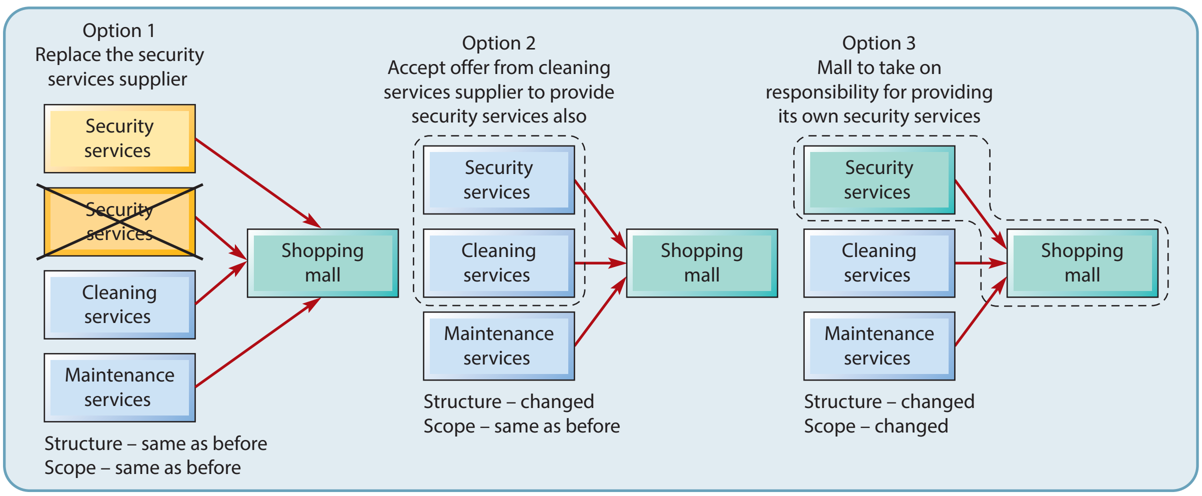

Figure 4.3 Three options for the shopping mall's supply network

DIAGNOSTIC QUESTION

How vertically integrated should the operation's network be?

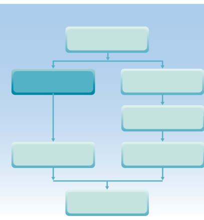

The scope of an operation's supply network determines the extent that an operation does things itself and the extent that it will rely on other operations to do things for it. This is often referred to as 'vertical integration' when it is ownership of whole operations that are being decided, or 'outsourcing' when individual activities are being considered. We will look at the 'outsourcing' decision in the next section. Vertical integration is the extent to which an organisation owns the network of which it is a part. It usually involves an organisation assessing the wisdom of acquiring suppliers or customers. And different companies, even in the same industry, can make very different decisions over how much and where in the network they want to be. Figure 4.4 illustrates the (simplified) supply network for the wind turbine power generation industry. Original equipment manufacturers (OEMs) assemble the wind turbine nacelle (the nacelle houses the generator and gearbox). Towers and blades are often built to the OEM's specifications, either in-house or by outside suppliers. Installing wind turbines involves assembling the nacelle, tower and blades on site, erecting the tower and connecting to the electricity network. The extent of vertical integration varies by company and component. The three companies illustrated in Figure 4.4 have all chosen different vertical integration strategies. Company A is primarily a nacelle designer and manufacturer that also make the parts. Company B is primarily an installer

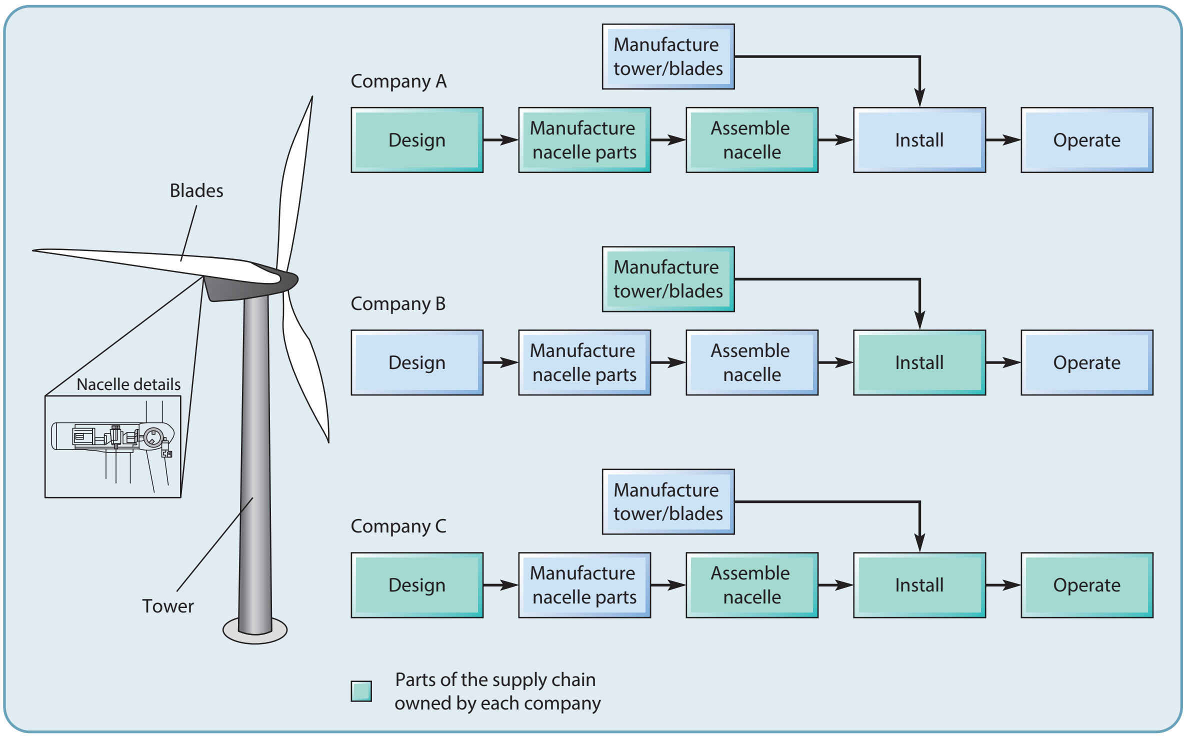

Figure 4.4 Three companies operating in the wind power generation industry with different vertical integration positions

that also makes the tower and blades (but buys-in the nacelle itself). Company C is primarily an operator that generates electricity and also designs and assembles the nacelles as well as installing the whole tower (but it outsources the manufacture of the nacelle parts, tower and blades).

Generally, an organisation's vertical integration strategy can be defined in terms of:

- *The direction of any integration* – Does it expand by buying one of its suppliers (backward or 'upstream' vertical integration) or should it expand by buying one of its customers (forward or 'downstream' vertical integration)? Backward vertical integration, by allowing an organisation to take control of its suppliers, is often used either to gain cost advantages or to prevent competitors gaining control of important suppliers. Forward vertical integration, on the other hand, takes an organisation closer to its markets and allows more freedom for an organisation to make contact directly with its customers, and possibly sell complementary products and services.
- *The extent of the process span of integration* – Some organisations deliberately choose not to integrate far, if at all, from their original part of the network. Alternatively some organisations choose to become very vertically integrated.
- *The balance among the vertically integrated stages* – This is not strictly about the ownership of the network. It refers to the amount of the capacity at each stage in the network that is devoted to supplying the next stage. So a totally balanced network relationship is one where one stage produces only for the next stage in the network and totally satisfies its requirements.

Advantages and disadvantages of vertical integration

The decision as to whether to vertically integrate in a particular set of circumstances is largely a matter of a business balancing the following advantages and disadvantages as they apply to them.

Perceived advantages of vertical integration

Although extensive vertical integration is no longer as popular as it once was, there are still companies who find it advantageous to own several sequential stages of their supply network. Indeed very few companies are anywhere close to 'virtual', with no vertical integration of stages whatsoever. What then are the reasons why companies still choose to vertically integrate? Most justifications for vertical integration fall under four categories. These are:

1. ***It secures dependable access to supply or markets*** – The most fundamental reasons for engaging in some vertical integration are that it can give a more secure supply or bring a business closer to its customers. One reason why the oil companies, who sell gasoline, are also engaged in extracting it is to ensure long-term supply. In some cases, there may not even be sufficient capacity in the supply market to satisfy the company. It therefore has little alternative but to supply itself. Downstream vertical integration can give a firm greater control over its market positioning. For example, Apple has always adopted a supply network model that integrates hardware and software with both its hardware and software designed by Apple.
2. ***It may reduce costs*** – The most common argument here is that 'We can do it cheaper than our supplier's price'. Such statements are often made by comparing the marginal direct cost incurred by a company in doing something itself against the price it is paying to buy the product or service from a supplier. But costs saving should also take into account start-up and learning costs. A more straightforward case can be made when there are technical advantages of integration. For example, producing aluminium kitchen foil

involves rolling it to the required thickness and then 'slitting' it into the finished widths. Performing both activities in-house saves the loading and unloading activity and the transportation to another operation. Vertical integration also reduces the 'transaction costs' of dealing with suppliers and customers. Transaction costs are expenses, other than price, which are incurred in the process of buying and selling, such as searching for and selecting suppliers, setting up monitoring arrangements, negotiating contracts, and so on. If transaction costs can be lowered to the point where the purchase price plus transaction costs is less than the internal cost, there is little justification for the vertical integration of the activity.

3. ***It may help to improve product or service quality*** – Sometimes vertical integration can be used to secure specialist or technological advantage by preventing product and service knowledge getting into the hands of competitors. The exact specialist advantage may be anything from the 'secret ingredient' in fizzy drinks through to a complex technological process. In either case the argument is the same. This process gives us the key identifying factor for our products and services. Vertical integration therefore is necessary to the survival of product or service uniqueness.
4. ***It helps in understanding other activities in the supply network*** – Some companies, even those who are famous for their rejection of traditional vertical integration, do choose to own some parts of the supply network other than what they regard as core. For example, McDonald's, the restaurant chain, although largely franchising its retail operations, does own some retail outlets. How else, it argues, could it understand its retail operations so well?

Perceived disadvantages of vertical integration

The arguments against vertical integration tend to cluster around a number of observed disadvantages of those companies that have practised vertical integration extensively. These are:

1. ***It creates an internal monopoly*** – Operations, it is argued, will only change when they see a pressing need to do so. Internal supply is less subject to the normal competitive forces that keep operations motivated to improve. If an external supplier serves its customers well, it will make higher profits; if not, it will suffer. Such incentives and sanctions do not apply to the same extent if the supplying operation is part of the same company.
2. ***You can't exploit economies of scale*** – Any activity that is vertically integrated within an organisation is probably also carried out elsewhere in the industry. But the effort it puts into the process will be a relatively small part of the sum total of that activity within the industry. Specialist suppliers who can serve more than one customer are likely to have volumes larger than any of their customers could achieve doing things for themselves. This allows specialist suppliers to reap some of the cost benefits of economies of scale, which can be passed on in terms of lower prices to their customers.
3. ***It results in loss of flexibility*** – Heavily vertically integrated companies, by definition, do most things themselves. This means that a high proportion of their costs will be fixed costs. They have, after all, invested heavily in the capacity that allows them to do most things in-house. A high level of fixed costs relative to variable costs means that any reduction in the total volume of activity can easily move the economics of the operation close to, or below, its break-even point.
4. ***It cuts you off from innovation*** – Vertical integration means investing in the processes and technologies necessary to produce products and services in-house. But, as soon as that investment is made the company has an inherent interest in maintaining it. Abandoning such investments can be both economically and emotionally difficult. The temptation is always to wait until any new technology is clearly established before admitting that one's

own is obsolete. This may lead to a tendency to lag in the adoption of new technologies and ideas.

5. ***It distracts you from core activities (loss of focus)*** – The final, and arguably most powerful, case against vertical integration concerns any organisation's ability to be technically competent at a very wide range of activities. All companies have things that it needs to be good at. And it is far easier to be exceptionally good at something if the company focuses exclusively on it, rather than being distracted by many other things. Vertical integration, by definition, means doing more things, which can distract from the (few) particularly important things.

DIAGNOSTIC QUESTION

How do operations decide what to do in-house and what to outsource?

Theoretically, the 'vertical integration' decision and the 'outsourcing' decision are almost the same thing. The difference between them is one of scale. Vertical integration is a term that is usually (but not always) applied to whole operations. So, buying a supplier because you want to deny their products to a competitor, or selling the part of your business that services your products to a specialist servicing company that can do the job better, is a vertical integration decision. Outsourcing usually applies to smaller sets of activities that have previously been performed in-house. Deciding to ask a specialist laboratory to perform some quality tests that your own quality control department used to do, or having your call (contact) centre taken over and run by a larger call-centre company, are both outsourcing decisions.

Outsourcing is also known as the 'do-or-buy' decision, and has become an important issue for most businesses. This is because, although most companies have always outsourced some of their activities, a larger proportion of direct activities are now bought from suppliers. Also many indirect and administrative processes are now outsourced. This is often referred to as business process outsourcing (BPO). Financial service companies, in particular, outsource some of their more routine back-office processes. In a similar way, many processes within the human resource function, from simple payroll services through to more complex training and development processes, are outsourced to specialist companies. The processes may still be physically located where they were before, but the outsourcing service provider manages the staff and technology. The reason for doing this is often primarily to reduce cost. However, there can sometimes also be significant gains in the quality and flexibility of service offered.

Making the outsourcing decision

Outsourcing is rarely a simple decision. Operations in different circumstances with different objectives are likely to make different decisions. Yet the question itself is relatively simple, even if the decision itself is not: 'Does in-house or outsourced supply in a particular set of circumstances give the appropriate performance objectives that it requires to compete more effectively in its markets?' For example, if the main performance objectives for an operation are dependable delivery and meeting short-term changes in customers' delivery requirements, the key question should be: 'How does in-house or outsourcing give better

Table 4.1 How in-house and outsourced supply may affect an operation's performance objectives

Performance objective	*'Do it yourself' In-house supply*	*'Buy it in' Outsourced supply*
Quality	The origins of any quality problems are usually easier to trace in-house and improvement can be more immediate but there can be some risk of complacency.	Supplier may have specialised knowledge and more experience, also may be motivated through market pressures, but communication more difficult.
Speed	Can mean synchronised schedules which speeds throughput of materials and information, but if the operation has external customers, internal customers may be low priority.	Speed of response can be built into the supply contract where commercial pressures will encourage good performance, but there may be significant transport/delivery delays.
Dependability	Easier communications can help dependability, but, if the operation also has external customers, internal customers may receive low priority.	Late delivery penalties in the supply contract can encourage good delivery performance, but organisational barriers may inhibit communication.
Flexibility	Closeness to the real needs of a business can alert the in-house operation to required changes, but the ability to respond may be limited by the scale and scope of internal operations.	Outsourced suppliers may be larger with wider capabilities than in-house suppliers and more ability to respond to changes, but may have to balance conflicting needs of different customers.
Cost	In-house operations do not have to make the margin required by outside suppliers so the business can capture the profits which would otherwise be given to the supplier, but relatively low volumes may mean that it is difficult to gain economies of scale or the benefits of process innovation.	Probably the main reason why outsourcing is so popular. Outsourced companies can achieve economies of scale and they are motivated to reduce their own costs because it directly impacts on their profits, but costs of communication and coordination with supplier need to be taken into account.

OPERATIONS PRINCIPLE

Assessing the advisability of outsourcing should include how it impacts on relevant performance objectives.

dependability and delivery flexibility performance?' This means judging two sets of opposing factors – those that give the potential to improve performance, and those that work against this potential being realised. Table 4.1 summarises some arguments for in-house supply and outsourcing in terms of each performance objective.

Incorporating strategic factors into the outsourcing decision

Although the effect of outsourcing on the operation's performance objective is important, there are other factors that companies take into account when deciding if outsourcing an activity is a sensible option. For example, if an activity has long-term strategic importance to a company, it is unlikely to outsource it. For example, a retailer might choose to keep the design and development of its website in-house because it plans to move into web-based retailing at some point in the future, even though specialists could perform the activity at less cost. Nor would a company usually outsource an activity for which it had specialised skills or knowledge. For example, a company making laser printers may have built up specialised knowledge in the production of sophisticated laser drives. This capability may allow it to introduce product or process innovations in the future. It would be foolish to 'give away' such capability. After these two more strategic factors have been considered, the company's operations performance can be taken into account. Obviously if its operation's performance is already too superior to any potential supplier, it would be unlikely to outsource the activity. But even if its performance was currently below that of potential suppliers, it may not outsource the activity if it feels that it could significantly improve its performance. Figure 4.5 illustrates this decision logic.

OPERATIONS PRINCIPLE

Assessing the advisability of outsourcing should include consideration of the strategic importance of the activity and the operation's relative performance.

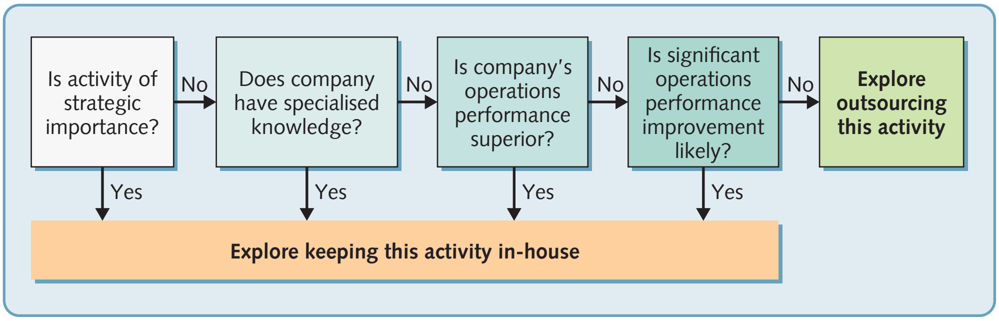

Figure 4.5 The decision logic of outsourcing

EXAMPLE

Samsung's subcontracted success

One of the best-known cautionary tales that illustrates the inherent dangers involved in subcontracting is that of how General Electric lost its microwave oven business. Although Japanese domestic appliance manufacturers, such as Matsushita and Sanyo, dominated the global microwave industry at the beginning of the 1980s General Electric (GE) was enjoying reasonable success in the US market with its purpose-designed microwave oven plant in Maryland. However, the company soon came under price pressures from Japanese competitors. What seemed an obvious solution was to subcontract the production of some of its more basic models, where margins were relatively small. GE explored the idea of subcontracting these models to one of its main rivals Matsushita, even though giving one of its main competitors such an advantage was considered risky. GE also found a small, but go-getting, Korean company who were already selling very simple (and very cheap) models in the US. GE decided to continue making the top of the range models itself, subcontract its cheaper models to Matsushita, but also place a small order of 15,000 units of its cheaper models to the Korean company, partly to see whether they could cope with the order. Of course it also made sense for GE to send their own engineers to help the Korean company and ensure that quality standards would be maintained. The GE engineers found that, although the Korean company had little knowledge, they were very willing to learn. Eventually, the Korean's production line started producing reasonable quality products, still at very low prices. Over time, the Korean company was given more and more orders by GE, who found that they were making more profit from the Korean sourced products than those coming out of their Maryland plant. This became particularly important as the market continued to mature and costs came under increased pressure. The Maryland plant attempted to cut its own costs but this proved especially difficult with so much of its volume now subcontracted to the Korean company. In the end the Maryland plant was closed and GE withdrew entirely from the microwave oven (indeed the whole domestic appliance) market. And the Korean company? It was called Samsung, and within 10 years of starting to make them, it became the world's largest manufacturer of microwave ovens.

Outsourcing and offshoring

Two supply network strategies that are often confused are those of outsourcing and offshoring. Outsourcing means deciding to buy-in products or services rather than perform the activities in-house. Offshoring means obtaining products and services from operations that are based outside one's own country. Of course, one may both outsource and offshore. Offshoring is very closely related to outsourcing and the motives for each may be similar. Offshoring to a lower cost region of the world is usually done to reduce an operation's overall costs, as is outsourcing to a supplier who has greater expertise or scale or both.

Globalisation, 're-shoring' and technology

Almost since the start of the industrial revolution businesses across the world have been forging ever-closer relationships. Especially in the last few decades, the use of geographically dispersed suppliers to outsource at least some activities has become routine. This has been driven partly by labour cost differentials, partly by cheap and efficient telecommunication between businesses, partly by trade deals, and partly by reducing transport costs. This is 'globalisation', where products, raw materials, money, technology and ideas move (relatively) smoothly across national boundaries. Apple and others can design their products in California, while Chinese assembly lines assemble them. A French aerospace company can direct the activities of its Brazilian suppliers almost as effectively as if they were in the next town. Bill Clinton, the ex-American President called globalisation 'the economic equivalent of a force of nature, like wind or water. It pushes countries to specialise and swap, making them richer, and the world smaller.' But increasingly some economists and business commentators question whether the boom in globalised operations is over. Some cite protectionist pressures in some developed countries. Others see rising wages in (previously) less developed countries as reducing cost differentials. In addition, the operations-related advantages of sourcing from nearby suppliers can be significant. Reducing reliance on complicated international supply chains can save transport and inventory costs, and is less polluting and potentially less prone to reputational risk if far-off suppliers misbehave. It also could increase supply flexibility. For example, the Spanish fast fashion brand, Zara, manufactures some of its 'steady selling' items in low-cost factories in Asia, but makes its garments with less predictable demand closer to its markets so that it can respond quickly to changing fashions. Developments in technology could reinforce this so-called 're-shoring' process. Automation may encourage a trend towards 'radical insourcing', where developed countries no longer need to outsource production to countries where wages are low. For example, in 2016 Adidas, the shoe manufacturer, announced that for the first time in 30 years it would bring a small proportion of its production back to Germany to be made in a highly automated factory in Bavaria.

DIAGNOSTIC QUESTION

What configuration should a supply network have?

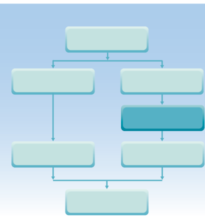

'Configuring' a supply network means determining its overall pattern. In other words, what should be the pattern, shape or arrangement of the various operations that make up the supply network? Even when an operation does not directly own, or even control, other operations in its network, it may still wish to change the shape of the network. This involves attempting to manage network behaviour by reconfiguring the network to change the nature of the relationships between them. Reconfiguring a supply network sometimes involves parts of the operation being merged – not necessarily in the sense of a change of ownership of any parts of an operation, but rather in the way responsibility is allocated for carrying out activities. The most common example of network reconfiguration has come through the many companies that have recently reduced the number of their direct suppliers. The complexity of dealing with many hundreds of suppliers may both be expensive for an operation and (sometimes more importantly) prevent the operation from developing a close relationship with a supplier. It is not easy to be close to hundreds of different suppliers.

Disintermediation

Another trend in some supply networks is that of companies within a network bypassing customers or suppliers to make contact directly with customers' customers or suppliers' suppliers. 'Cutting out the middle men' in this way is called disintermediation. An obvious example of this is the way the internet has allowed some suppliers to 'disintermediate' traditional retailers in supplying goods and services to consumers. So, for example, many services in the travel industry that used to be sold through retail outlets (travel agents) are now also available direct from the suppliers. The option of purchasing the individual components of a vacation through the websites of the airline, hotel, car-hire operation, and so on, is now easier for consumers. Of course, they may still wish to purchase an 'assembled' product from retail travel agents, which can have the advantage of convenience. Nevertheless, the process of disintermediation has developed new linkages in the supply network.

Co-opetition

One approach to thinking about supply networks sees any business as being surrounded by four types of players: suppliers, customers, competitors and complementors. Complementors enable one's products or services to be valued more by customers because they also can have the complementor's products or services, as opposed to when they have yours alone. Competitors are the opposite; they make customers value your product or service less when they can have their product or service, rather than yours alone. Competitors can also be complementors and vice versa. For example, adjacent restaurants may see themselves as competitors for customers' business. A customer standing outside and wanting a meal will choose between the two of them. Yet, in another way they are complementors. Would that customer have come to this part of town unless there was more than one restaurant to choose from? Restaurants, theatres, art galleries and tourist attractions generally, all cluster together in a form of cooperation to increase the total size of their joint market. It is important to distinguish between the way companies cooperate in increasing the total size of a market and the way in which they then compete for a share of that market. Customers and suppliers, it is argued, should have 'symmetric' roles. Harnessing the value of suppliers is just as important as listening to the needs of customers. Destroying value in a supplier in order to create it in a customer does not increase the value of the network as a whole. So, pressurising suppliers will not necessarily add value. In the long term it creates value for the total network to find ways of increasing value for suppliers 'as well as customers. All the players in the network, whether they are customers, suppliers, competitors or complementors, can be both friends and enemies at different times. The term used to capture this idea is 'co-opetition'.

The idea of the 'business ecosystem'[4]

An idea that is closely related to that of co-opetition in supply networks is that of the 'business ecosystem'. It can be defined as: *'An economic community supported by a foundation of interacting organizations and individuals—the organisms of the business world. The economic community produces goods and services of value to customers, who are themselves members of the ecosystem. The member organisms also include suppliers, lead producers, competitors, and other stakeholders. Over time, they coevolve their capabilities and roles, and tend to align themselves with the directions set by one or more central companies'.* [5] One of the main differences between this idea and that of the supply network generally is the inclusion in the idea of the ecosystem of businesses that may have no or little direct relationship with the main supply network, yet exist only because of that network. They interact with each other, predominantly complementing or contributing significant components of the value proposition for customers. Many examples come from the technology industries. The innovative products and services that are developed in the technology sectors cannot evolve in a vacuum. They need to attract a whole range of resources, drawing in expertise, capital, suppliers and customers to create cooperative networks. For example, the app developers that develop applications for particular operating system platforms may not be 'suppliers', as such, but the relationship between them and the supply network that supplies the mobile device is mutually beneficial. Building an ecosystem of developers around a core product can increase its value to the end customer and by doing so complements the usage of the core product. Such an ecosystem of complementary products and services can also create significant barriers to entry for new competitors. Any possible competitors would not only have to compete with the core product, but they would also have to compete against the entire ecosystem of complementary products and services.

The terminology and metaphors used to describe business ecosystems are obviously based on that used to describe 'natural' biological systems, where elements in the 'ecosystem' affect and are affected by the others. This creates a constantly evolving set of relationships where, if they are to survive, businesses must be flexible, adaptable and preferably innovative. For an ecosystem to thrive, the relationships between elements (businesses in this case) must communicate, establish trust, share information, collaborate, experiment and develop in a mutually supportive symbiotic manner. The comparison with the natural biological ecosystem is also important because it emphasises that the relationships between things matter and that, to some extent, everything in a supply network touches everything else.

Describing supply networks – dyads and triads

The supply network that was illustrated in Figure 4.2 is, of course, a simplification. Any realistic supply network diagram will be much more complex. There are many operations, all interacting in different ways, to produce end products and services. Because of this, and to understand them better, supply network academics and professionals often choose to focus on the individual interaction between two specific operations in the network. This is called a 'dyadic' (simply meaning 'two') interaction, or dyadic relationship, and the two operations are referred to as a 'dyad'. So if one wanted to examine the interactions that a focal operation had with one of its suppliers and one of its customers, one would examine the two dyads of 'supplier – focal operation', and 'focal operation – customer', see Figure 4.6(a). For many years most discussion (and research) on supply networks was based on dyadic relationships. This is not surprising as all relationships in a network are based on the simple dyad. However, more recently, and certainly when examining service supply networks, many authorities make the point that dyads do not reflect the real essence of a supply network. Rather, they say, it is triads, not dyads, that are the basic elements of a supply network, see Figure 4.6(b). No matter how complex a network, it can be broken down into a collection of triadic interactions. The idea of triads is especially relevant in service supply networks. Operations are increasingly outsourcing the delivery of some aspects of their service to specialist providers, who deal directly with customers on behalf of the focal operation (more usually called the 'buying operation', or just 'buyer' in this context). For example, Figure 4.6(b) illustrates the common example of an airline contracting a specialist baggage handling company to provide services to its customers on its behalf. Similarly, internal services are increasingly outsourced to form internal triadic relationships. For example, if a company outsources its IT operations, it is forming a triad between whoever is purchasing the service on behalf of the company, the IT service provider, and the employees who use the IT services.

Thinking about supply networks as a collection of triads rather than dyads is strategically important. First, it emphasises the dependence that organisations place on their suppliers' performance when they outsource service delivery. A supplier's service performance makes up an important part of how the buyer's performance is viewed. Second, the control that the buyer of the service has over service delivery to its customer is diminished in a triadic relationship. In a conventional supply chain, with a series of dyadic relationships, there is the opportunity to intervene before the customer receives the product or service. However, products, or services in triadic relationships bypass the buying organisation and go directly from provider to customer. Third, and partially as a consequence of the previous point, in triadic relationships

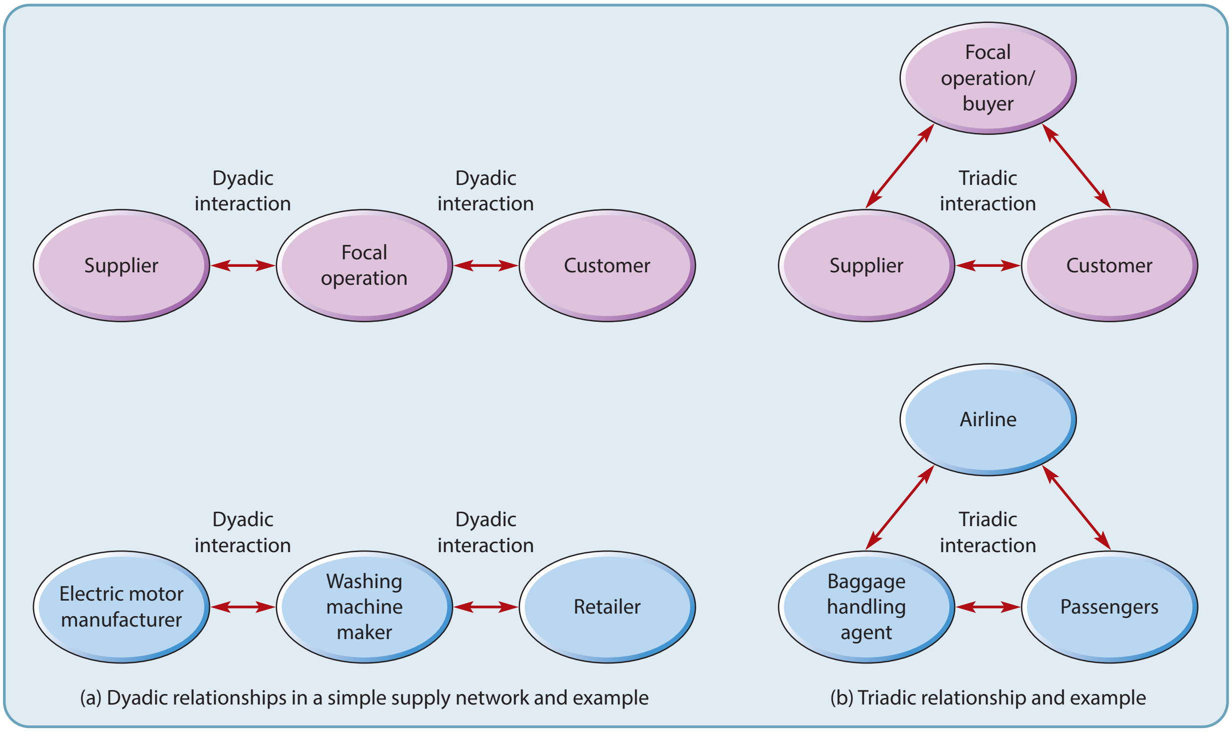

Figure 4.6 **Dyadic and triadic relationships in two simple supply networks and examples**

the direct link between service provider and customer can result in power gradually transferring over time from the buying organisation to the supplier that provides the service. Fourth, it becomes increasingly difficult for the buying organisation to understand what is happening between the supplier and customer at a day-to-day level. It may not even be in the supplier's interests to be totally honest in giving performance feedback to the buyer. Finally, this closeness between supplier and customer, if it excludes the buyer, could prevent the buyer from building important knowledge. For example, suppose a specialist equipment manufacturer has outsourced the maintenance of its equipment to a specialist provider of maintenance services. The ability of the equipment manufacturer to understand how its customers are using the equipment, how the equipment is performing under various conditions, how customers would like to see the equipment improved, and so on, is lost. The equipment manufacturer may have outsourced the cost and trouble of providing maintenance services, but it has also outsourced the benefits and learning that come from direct interaction with customers.

DIAGNOSTIC QUESTION

How much capacity should operations plan to have?

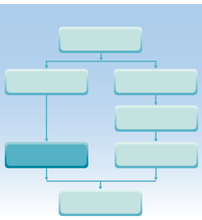

The next set of 'structure' decisions concern the size or capacity of each part of the supply network. Here we shall treat capacity in a general long-term sense. The specific issues involved in measuring and adjusting capacity in the medium and short terms are examined in Chapter 8.

The optimum capacity level

Most organisations have decisions to make about how big (in terms of capacity) they want to be. A chain of truck service centres, for example, might operate centres that have various capacities. The effective cost of running each centre will depend on the average service bay occupancy. Low occupancy, because of few customers, will result in a high cost per customer served because the fixed costs of the operation are being shared between few customers. As demand, and therefore service bay occupancy, increases the cost per customer will reduce. However, operating at very high levels of capacity utilisation (occupancy levels close to capacity) can mean longer customer waiting times and reduce customer service. There may also be less obvious cost penalties of operating centres at levels close to its nominal capacity. For example, long periods of overtime may reduce productivity levels, as well as costing more in extra payments to staff; utilising bays at very high levels reduces the time available for maintenance and cleaning which may increase the number of breakdowns, reduce effective life, and so on. This usually means that average costs start to increase after a point that will often be lower than the theoretical capacity of the operation.

OPERATIONS PRINCIPLE

All types of operation exhibit economy of scale effects where operating costs reduce as the scale of capacity increases.

The blue curves in Figure 4.7 show this effect for the service centres of 5-, 10- and 15-bay capacity. As the nominal capacity of the centres increases, the lowest cost point at first reduces. This is because the fixed costs of any operation do not increase proportionately as its capacity increases. A 10-bay centre has less than twice the fixed costs of a 5-bay centre. The capital costs of constructing the operations do not increase proportionately to their capacity. A 10-bay centre costs less to build than twice the cost of building a 5-bay centre. These two factors, taken together, are often referred to as economies of scale – a universal concept that applies (up to a point) to all types of operation. However, economies of scale do not go on forever. Above a certain size, the lowest cost point on curves such as that shown in Figure 4.7 may increase. This occurs because of what are called diseconomies of scale, two of which are particularly important. First, complexity costs increase as size increases. The communications and coordination effort necessary to manage an operation tends to increase faster than capacity. Although not seen as a direct cost, this can nevertheless be very significant. Second, a larger centre is more likely to be partially underutilised because demand within a fixed location will be limited. The equivalent in operations that

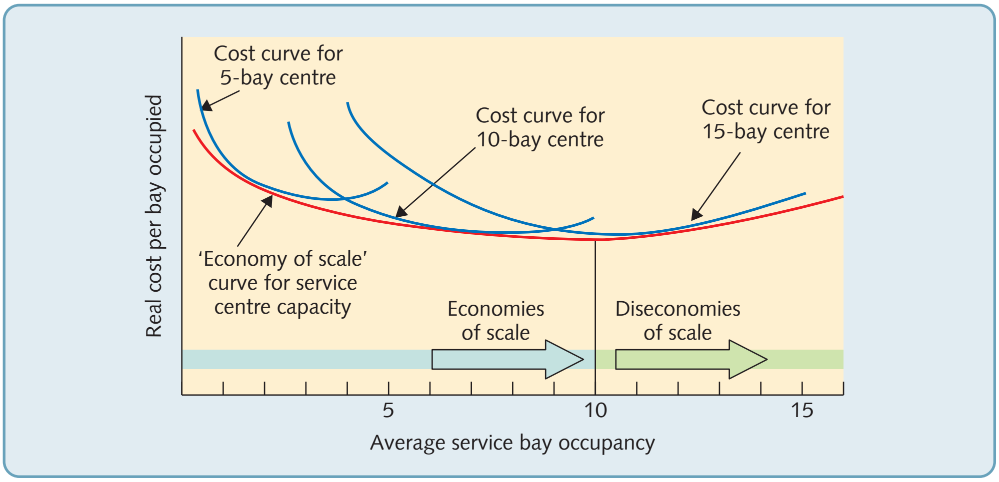

Figure 4.7 Unit cost curves for individual truck service centres of varying capacities

OPERATIONS PRINCIPLE

Diseconomies of scale increase operating costs above a certain level of capacity resulting in a minimum cost level of capacity.

process physical items is transportation costs. For example, if a manufacturer supplies the whole of its European market from one major plant in Denmark, all supplies may have to be brought in from several countries to the single plant and all products shipped from there throughout Europe.

Being small may have advantages

Although large-scale capacity operations will usually have a cost advantage over smaller units, there are also potentially significant advantages that can be exploited by small-scale operations. One significant research study showed that small-scale operations can provide significant advantages in the following four areas.[6]

1. They allow businesses to locate near to 'hot spots' that can tap into local knowledge networks. Often larger companies centralise their research and development efforts, so losing touch with where innovative ideas are generated.
2. Responding rapidly to regional customer needs and trends by basing more and smaller units of capacity close to local markets.
3. Taking advantage of the potential for human resource development, by allowing staff a greater degree of local autonomy. Larger scale operations often have longer career paths with fewer opportunities for 'taking charge'.
4. Exploring radically new technologies by acting in the same way as a smaller more entrepreneurial rival. Larger, more centralised development activities are often more bureaucratic than smaller scale agile centres of development.

EXAMPLE

Economies of scale in heart surgery and shipping[7]

Don't think that the idea of economies of scale only applies to manufacturing operations. It is a universal concept. Here are just two examples.

In the 1,000-bed Narayana Hrudayalaya hospital, in Bangalore India, Dr Devi Shetty has created what, according to *Forbes* magazine, is the world's largest heart factory. It is a radical new approach, he said, and proves that economies of scale can transform the cost of cardiology. The hospital has 42 surgeons who perform 6,000 heart operations each year, including 3,000 on children. This makes the hospital the busiest facility of its type in the world. And it's needed; it is estimated that India requires 2.5 million heart operations every year, yet only 90,000 are performed. '*It's a numbers game,*' said Dr Shetty, who has performed 15,000 heart operations. '*Surgeons are technicians. The more practice they get, the more specialised they become and the better the results.*' The result is that costs are slashed and the hospital can be profitable even though many patients are poor. The hospital's charges for open-heart surgery are, on average, a tenth of the cost of the cheapest procedures in the United States. But even then, treatment is too expensive for many, so wealthier patients are charged more to subsidise the poorest.

The *Eleonora Maersk* is one of seven ships in her class that are owned by Maersk Lines, the world's biggest container-shipping company. They are among the biggest ships ever built: almost 400 m long (the length of four football pitches). The *Eleonora Maersk* is also powerful; it has the largest internal-combustion engine ever built, as powerful as 1,000 family cars, which enables it to move all its cargo from China to Europe in just over 3 weeks. Yet the ship is so automated that it requires only 13 people to crew it. On board, the ship can carry 15,000 20-foot containers, each one of which can hold 70,000 T-shirts. It is these economies of scale that allow a T-shirt made in China to be sent to the Netherlands for just 2.5 cents. And the economies of scale involved in building and running these ships means that things will get bigger still. Hoping to drive costs down further, the ship's owners have ordered 20 even larger ships with a capacity of 18,000 twenty-foot containers, costing $200 million each.

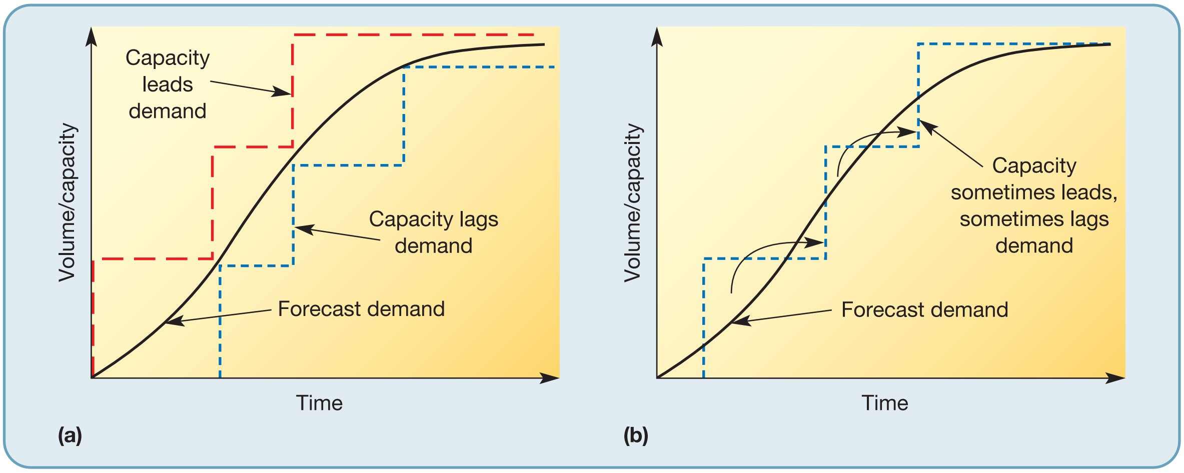

Figure 4.8 **(a) Capacity-leading and capacity-lagging strategies, (b) Smoothing with inventories means using the excess capacity in one period to produce inventory that supplies the under-capacity period**

The timing of capacity change

Changing the capacity of any operation in a supply network is not just a matter of deciding on its optimum capacity. The operation also needs to decide when to bring 'on-stream' new capacity. For example, Figure 4.8 shows the forecast demand for a manufacturer's new product. In deciding *when* new capacity is to be introduced the company can mix three strategies.

1. Capacity is introduced to generally lead demand – timing the introduction of capacity in such a way that there is always sufficient capacity to meet forecast demand.
2. Capacity is introduced to generally lag demand – timing the introduction of capacity so that demand is always equal to or greater than capacity.
3. Capacity is introduced to sometimes lead and sometimes lag demand, but inventory built up during the 'lead' times is used to help meet demand during the 'lag' times. This is called 'smoothing with inventory'.

Each strategy has its own advantages and disadvantages. These are shown in Table 4.2. The actual approach taken by any company will depend on how it views these advantages and

Table 4.2 **The arguments for and against pure leading, pure lagging and smoothing with inventory strategies of capacity timing**

Advantages	*Disadvantages*
Capacity-leading strategies	
Always sufficient capacity to meet demand,therefore revenue is maximised and customers satisfied	Utilisation of the plants is always relatively low, therefore costs will be high
Most of the time there is a 'capacity cushion' that can absorb extra demand if forecasts are pessimistic	Risks of even greater (or even permanent) over-capacity if demand does not reach forecast levels
Any critical start-up problems with new operations are less likely to affect supply	Capital spending on capacity will be early
Capacity-lagging strategies	
Always sufficient demand to keep the operation working at full capacity, therefore unit costs are minimised	Insufficient capacity to meet demand fully, therefore reduced revenue and dissatisfied customers
Over-capacity problems are minimised if forecasts prove optimistic	No ability to exploit short-term increases in demand

(Continued)

Table 4.2 The arguments for and against pure leading, pure lagging and smoothing with inventory strategies of capacity timing (*Continued*)

Advantages	*Disadvantages*
Capital spending on the operation is delayed	Under-supply position even worse if there are start-up problems with the new operations
Smoothing with inventory strategies	
All demand is satisfied, therefore customers are satisfied and revenue is maximised	The cost of inventories in terms of working capital requirements can be high. This is especially serious at a time when the company requires funds for its capital expansion
Utilisation of capacity is high and therefore costs are low	Risks of product deterioration and obsolescence
Very short-term surges in demand can be met from inventories	

OPERATIONS PRINCIPLE

Capacity-leading strategies increase opportunities to meet demand. Capacity-lagging strategies increase capacity utilisation.

OPERATIONS PRINCIPLE

Using inventories to overcome demand–capacity imbalance tends to increase working capital requirements.

disadvantages. For example, if the company's access to funds for capital expenditure is limited, it is likely to find the delayed capital expenditure requirement of the capacity-lagging strategy relatively attractive. Of course, the third strategy, 'smoothing with inventory', is only appropriate for operations that produce products that can be stored. Customer-processing operations such as a hotel cannot satisfy demand in one year by using rooms that were vacant the previous year.

Break-even analysis of capacity expansion

An alternative view of capacity expansion can be gained by examining the cost implications of adding increments of capacity on a break-even basis. Figure 4.9 shows how increasing capacity can move an operation from profitability to loss. Each additional unit of capacity results in a *fixed-cost break* that is a further lump of expenditure that will have to be incurred before any further activity can be undertaken in the operation. The operation is unlikely to be profitable at very low levels of output. Eventually, assuming that prices are greater than marginal costs, revenue will exceed total costs. However, the level of profitability at the point where the output level is equal to the capacity of the operation may not be sufficient to absorb all the extra fixed costs of a further increment in capacity. This could make the operation unprofitable in some stages of its expansion.

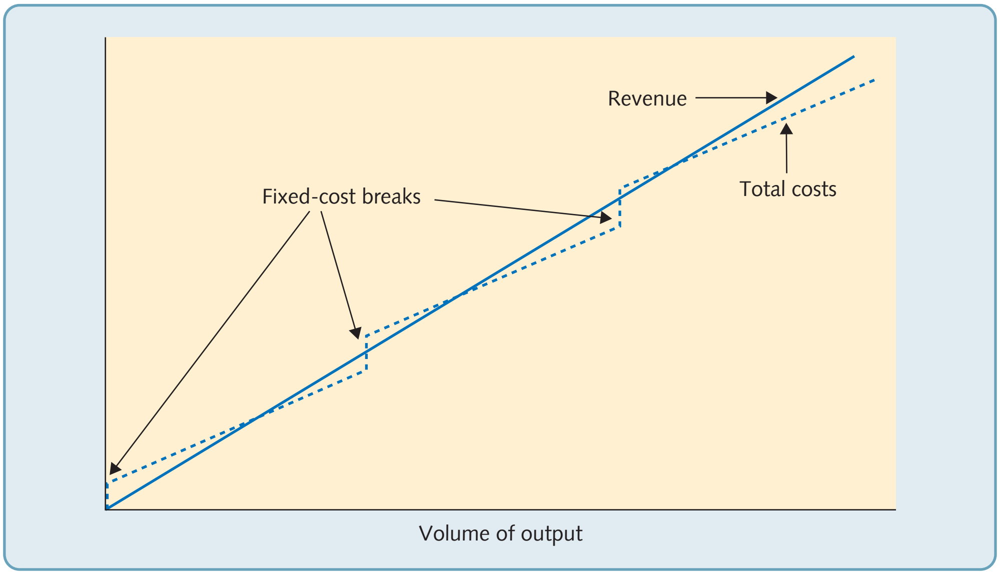

Figure 4.9 Repeated incurring of fixed costs can raise total costs above revenue

WORKED EXAMPLE

A specialist graphics company is investing in a new machine which enables it to make high-quality prints for its clients. Demand for these prints is forecast to be around 100,000 units in year 1 and 220,000 units in year 2. The maximum capacity of each machine the company will buy to process these prints is 100,000 units per year. They have a fixed cost of €200,000 per year and a variable cost of processing of €1 per unit. The company believe they will be able to charge €4 per unit for producing the prints. What profit are they likely to make in the first and second years?

$$\text{Year 1 demand} = 100{,}000 \text{ units; therefore the company will need one machine}$$

$$\begin{aligned}
\text{Cost of producing prints} &= \text{fixed cost for one machine} + \text{variable cost} \times 100{,}000\\
&= €200{,}000 + (€1 \times 100{,}000)\\
&= €300{,}000\\
\text{Revenue} &= \text{demand} \times \text{price}\\
&= 100{,}000 \times €4\\
&= €400{,}000\\
\text{Therefore profit} &= €400{,}000 - €300{,}000\\
&= €100{,}000
\end{aligned}$$

$$\text{Year 2 demand} = 220{,}000\text{; therefore the company will need three machines}$$

$$\begin{aligned}
\text{Cost of manufacturing} &= \text{fixed cost for three machines} + \text{variable cost} \times 220{,}000\\
&= (3 \times €200{,}000) + (€1 \times 220{,}000)\\
&= €820{,}000\\
\text{Revenue} &= \text{demand} \times \text{price}\\
&= 220{,}000 \times €4\\
&= €880{,}000\\
\text{Therefore profit} &= €880{,}000 - €820{,}000\\
&= €60{,}000
\end{aligned}$$

Note: the profit in the second year will be lower because of the extra fixed costs associated with the investment in the two extra machines.

DIAGNOSTIC QUESTION

Where should operations be located?

The location of each operation in a supply network is both a key element in defining its structure, and also will have an impact on how the network operates in practice. Poor location of any operation in a supply network can have a significant impact, not just on their profits, but also those of others in the network. For example, siting a data centre where potential staff with appropriate skills will not live will affect its performance, and the service it gives its customers. Location decisions will usually have an effect on an operation's costs as well as its ability to serve its customers (and therefore its revenues). In addition, location decisions, once taken, are difficult to undo. The costs of moving an operation can be hugely expensive and the risks of inconveniencing customers very high. No operation wants to move very often.

Why relocate?

Not all operations can logically justify their location. Some are where they are for historical reasons. Yet, even the operations that are 'there because they're there' are implicitly making a decision not to move. Presumably, their assumption is that the cost and disruption involved in changing location would outweigh any potential benefits of a new location. When operations do move, it is usually for one or both of two reasons – changes in demand or changes in supply.

Changes in demand

If customer demand shifts it may prompt a change in location. For example, as garment manufacturers moved to Asia, suppliers of zips, threads, and so on, started to follow them. Changes in the volume of demand can also prompt relocation. To meet higher demand, an operation could expand its existing site, or choose a larger site in another location, or keep its existing location and find a second location for an additional operation; the last two options will involve a location decision. High-visibility operations may not have the choice of expanding on the same site to meet rising demand. A dry cleaning service may attract only marginally more business by expanding an existing site because it offers a local, and therefore convenient, service. Finding a new location for an additional operation is probably its only option for expansion.

Changes in supply

The other stimulus for relocation is changes in the cost, or availability, of the supply of inputs to the operation. For example, a mining or oil company will need to relocate as the minerals it is extracting become depleted. The reason why so many software companies located in India was the availability of talented, well-educated, but relatively cheap staff.

EXAMPLE

Rolls Royce in Singapore

When Rolls Royce, the aero engine manufacturer chose Seletar in Singapore to host its almost £400 million Asian expansion, it did so for a number of reasons.[8] First, the location had access to the skills and infrastructure to support technically complex manufacturing. Second, the company was no stranger to the region; it already serviced Singapore Airlines' engines at a special plant near Changi airport. Third, Asia is where the demand is. The world's fastest-growing airlines are in China, Singapore, Indonesia, India and in the Gulf. Fourth, the generous tax incentives offered by the Singaporean government played a part, as did the construction of a road from Seletar to Changi airport so that engines can be loaded on to the cargo planes that fly them to Rolls Royce's customers in Toulouse and Seattle. Yet, the company said that although important, these incentives were not as beneficial as the 'soft' factors that make Singapore so attractive. In particular, the City State's universities and colleges, which produce the skilled scientists, engineers and staff who are vital to producing products that cannot be allowed to fail. According to Jonathan Asherson, *'we think that the focus in Asia, from an education and training perspective, will continue to be in areas of technology and engineering. The talent pipeline that we need as an industry and company will remain solid. That will influence the thinking around our investments. You need to develop technologies and business models that adapt to increasing pressure on costs,*

increasing pressure on reliability and the environment. We've worked with government agencies around developing work skills, qualifications, and developing curricula for the polytechnics and universities, where we work with them to predict the requirement and work on how that pipeline of talent can be built. Singapore is quite flexible and nimble where they see the high multiplier effect of, for example, high-value-added manufacturing.'

Evaluating potential changes in location

Evaluating possible locations is almost always a complex task because the number of location options, the criteria against which they could be evaluated and the comparative rarity of a single location that clearly dominates all others, make the decision strategically sensitive. Furthermore, the decision often involves high levels of uncertainty. Neither the relocation activity itself, nor the operating characteristics of the new site could be as assumed when the decision was originally made. Because of this, it is useful to be systematic in terms of (a) identifying alternative options, and (b) evaluating each option against a set of rational criteria.

OPERATIONS PRINCIPLE

An operation should only change its location if the benefits of moving outweigh the costs of operating in the new location plus the cost of the move itself.

Identify alternative location options

The first relocation option to consider is not to relocate. Sometimes relocation is inevitable, but often staying put is a viable option. Even if seeking a new location seems the obvious way forward, it is worth evaluating the 'do nothing' option, if only to provide a 'base case' against which to compare other options. But in addition to the 'do nothing' option there should be a number of alternative location options. It is a mistake to consider only one location, but seeking out possible locations can be a time-consuming activity. Increasingly, for larger companies, the whole world offers possible locations. The implication of the globalisation of the location decision has been to increase both the number of options and the degree of uncertainty in their relative merits. The sheer number of possibilities makes the location decision impossible to 'optimise'. Rather, the process of identifying location options usually involves selecting a limited number of sites that represent different attributes. For example, a distribution centre, while always needing to be close to transport links, could be located in any of several regions and could either be close to population centres, or in a more rural location. The options may be chosen to reflect a range of both these factors. However, this assumes that the 'supply' of location options is relatively large, which is not always the case. In many retail location decisions, there are a limited number of High Street locations that become available at any point in time. Often, a retailer will wait until a feasible location becomes available and then decide whether to either take up that option or wait and take the chance that a better location becomes available soon. In effect, the location decision here is a sequence of 'take or wait' decisions.

Set location evaluation criteria

Although the criteria against which alternative locations can be evaluated will depend on circumstances, the following five broad categories are typical:

1. ***Capital requirements*** – The capital or leasing cost of a site is usually a significant factor. This will probably be a function of the location of the site and its characteristics. For example, the shape of the site and its soil composition can limit the nature of any buildings erected there. Access to the site is also likely to be important, as are the availability of utilities, and so on. In addition, the cost of the move itself may depend on which site is eventually chosen.
2. ***Market factors*** – Location can affect how the market, either in general, or as individual customers, perceive an operation. Locating a general hospital in the middle of the countryside may have many advantages for its staff, but it clearly would be very inconvenient for

its customers. Likewise, restaurants, stores, banks, petrol filling stations and many other high-visibility operations, must all evaluate how alternative locations will determine their image and the level of service they can give. The same arguments apply to labour markets. Location may affect the attractiveness of the operation in terms of staff recruitment and retention. For example, 'science parks' are usually located close to universities because they hope to attract companies who are interested in using the skills available at the university. But not all locations necessarily have appropriate skills available immediately. Staff at a remote call centre in the western islands of Scotland, who were used to a calm and tranquil life, became stunned by the aggressive nature of many callers to the call centre, some being reduced to tears by bullying customers. They had to be given assertiveness training by the call centre management.

3. ***Cost factors*** – Two major categories of cost are affected by location. The first is the costs of producing products or services. For example, labour costs can vary between different areas in any country, but are likely to be a far more significant factor when international comparisons are made; they can exert a major influence on the location decision, especially in some industries such as clothing, for example, where labour costs as a proportion of total costs are relatively high. Other cost factors, known as community factors, derive from the social, political and economic environment of its site. These include such factors as local tax rates, capital movement restrictions, government financial assistance, political stability, local attitudes to 'inward investment', language, local amenities (schools, theatres, shops, etc.), the availability of support services, the history of labour relations and behaviour, environmental restrictions and planning procedures. The second category of costs relates to both the cost of transporting inputs from their source to the location of the operation and the cost of transporting products and services from the location to customers. Whereas almost all operations are concerned to some extent with the former, not all operations are concerned with the latter, either because customers come to them (for example, hotels), or because their services can be 'transported' at virtually no cost (for example, some technology help desks). For supply networks that process physical items, however, transportation costs can be very significant.
4. ***Future flexibility*** – Because operations rarely change their location, any new location must be capable of being acceptable, not only under current circumstances, but also under possible future circumstances. The problem is that no one knows exactly what the future holds. Nevertheless, especially in uncertain environments, any evaluation of alternative locations should include some kind of scenario planning that considers the robustness of each in coping with a range of possible futures. Two types of flexibility could be evaluated for any location. The most common one is to consider the potential of the location for expansion to cope with increased activity levels. The second is the ability to adapt to changes in input or output factors. For example, suppliers or customers may themselves relocate in the future. If so, could the location still operate economically?
5. ***Risk factors*** – Closely related to the concept of future flexibility, is the idea of evaluating the risk factors associated with possible locations. Again, the risk criterion can be divided into 'transition risk' and 'long-term risk'. Transition risk is simply the risk that something may go wrong during the relocation process. Some possible locations might be intrinsically more difficult to move to than others. For example, moving to an already congested location could pose higher risks to being able to move as planned than moving to a more accessible location. Long-term risks could again include damaging changes in input factors such as exchange rates or labour costs, but can also include more fundamental security risks to staff or property.

Critical commentary

Probably the most controversial issue in supply network design is that of outsourcing. In many instances, there has been fierce opposition to companies outsourcing some of their processes. Trade Unions often point out that the only reason that outsourcing companies can do the job at lower cost is that they either reduce salaries, reduce working conditions, or both. Furthermore, they say, flexibility is only achieved by reducing job security. Employees who were once part of a large and secure corporation could find themselves as far less secure employees of a less benevolent employer with a philosophy of permanent cost cutting. Even some proponents of outsourcing are quick to point out the problems. There can be significant obstacles, including understandable resistance from staff who find themselves 'outsourced'. Some companies have also been guilty of 'outsourcing a problem'. In other words, having failed to manage a process well themselves, they ship it out rather than face up to why the process was problematic in the first place. There is also evidence that, although long-term costs can be brought down when a process is outsourced, there may be an initial period when costs rise as both sides learn how to manage the new arrangement.

- The idea of widening the discussion of supply networks to include the 'business ecosystem' concept, described earlier, is also not without its critics. Some see it as simply another management 'buzzword', indistinguishable from the longer established idea of the supply network. Other critics, who believe that the ecosystem metaphor is just a way for business to appear 'Green', have criticised the use of the term 'business ecosystem' by commentators and firms. They claim that the metaphor is used to suggest that the commercial relationships, on which almost all supply networks are based, have developed and are run using 'natural' values and therefore should be left to operate free from societal or government interference.

SUMMARY CHECKLIST

- ☐ Is the operation fully aware of all its first and second-tier suppliers' and customers' capabilities and requirements?
- ☐ Are the capabilities of suppliers and requirements of customers understood in terms of all aspects of operations performance?
- ☐ Does the operation have a view on how it would like to see its supply network develop over time, both in terms of scope and structure?
- ☐ Have the benefits of reducing the number of individual suppliers been explored?
- ☐ Are any parts of the supply network likely to become disintermediated, and have the implications of this been considered?
- ☐ Does the operation have an approach to how its treats others in the supply network who might be both complementors and competitors?
- ☐ Is the vertical integration/outsourcing issue always under review for possible benefits?
- ☐ Is outsourcing (or bringing back in-house) evaluated in terms of all the operation's performance objectives?
- ☐ Is there a rational set of criteria used for deciding whether (or not) to outsource?
- ☐ Is the optimum economy of scale for the different types of operation with the business periodically assessed?
- ☐ Are the various strategies for timing changes in capacity always evaluated in terms of their advantages and disadvantages?
- ☐ Are the fixed-cost breaks of capacity increase understood, and are they taken into account when increasing or decreasing capacity?
- ☐ Is the relocation decision ever considered?
- ☐ Have factors such as changes in demand or supply that may prompt relocation been considered?
- ☐ If considering relocation, are alternative locations always evaluated against each other and against a 'do nothing' option?
- ☐ Are sufficient location options being considered?
- ☐ Do location evaluation criteria include capital, market, cost, flexibility, and risk factors?

CASE STUDY

Aarens Electronic

Just outside Rotterdam in the Netherlands, Frank Jansen, the Chief Operating Officer of Aarens Electronic (AE) was justifiably proud of what he described as, . . . '*the most advanced machine of its type in the world, which will enable us to achieve new standards of excellence for our products requiring absolute cleanliness and precision*' . . . and . . . '*a quantum leap in harnessing economies of scale, new technology to provide* ***the*** *most advanced operation for years to come.*' The Rotterdam operation was joining AE's two existing operations in the Netherlands. They offered precision custom coating and laminating services to a wide range of customers, among the most important being, Phanchem to whom it supplied dry photoresist imaging films, a critical step in the manufacturing of microchips. Phanchem then processed the film further and sold it direct to microchip manufacturers

The Rotterdam operation

The decision to build the Rotterdam operation had been taken because the company believed that a new low-cost operation using 'ultra-clean' controlled environment technology could secure a very large part of Phanchem's future business – perhaps even an exclusive agreement to supply 100 per cent of their needs. When planning the new operation, three options were presented to AE's Executive Committee:

1 Expand an existing site by building a new machine within existing site boundaries. This would provide around 12 to 13 million square metres (MSM) per year of additional capacity and require around €19 million in capital expenditure.
2 Build a new facility alongside the existing plant. This new facility could accommodate additional capacity of around 15 MSM per year but, unlike option A, would also allow for future expansion. Initially, this would require around €22 million of capital.
3 Set up a totally new site with a much larger increment of capacity (probably around 25 MSM per year). This option would be more expensive; at least €30 million.

Frank Jansen and his team initially favoured option B but in discussion with the AE Executive Committee, opinion shifted towards the more radical option C. '*It may have been the highest risk option but it held considerable potential and it fitted with the AE Group philosophy of getting into high-tech specialised areas of business. So we went for it.*' (Frank Jansen) The option of a very large, ultra-clean, state-of-the-art facility also had a further advantage: it could change the economics of the photoresist imaging industry. In fact, global demand and existing capacity did not immediately justify investing in such a large increase in capacity. There was probably some overcapacity in the industry. But a large-capacity, ultra-clean type operation could provide a level of quality at such low costs that, if there were over-capacity in the industry, it would not be AE's capacity that would by lying idle.

Designing the new operation

During discussions on the design of the new operation, it became clear that there was one issue that was underlying all the team's discussions – how flexible should the process be? Should the team assume that they were designing an operation that would be dedicated exclusively to the manufacture of photoresist imaging film, and ruthlessly cut out any technological options that would enable it to manufacture other products, or should they design a more general-purpose operation that was suitable for photoresist imaging film, but could also make other products? It proved a difficult decision. The advantages of the more flexible option were obvious. '*At least it would mean that there was no chance of me being stuck with an operation and no market for it to serve in a couple of years' time.*' (Frank Jansen) But the advantages of a totally dedicated operation were less obvious, although there was a general agreement that both costs and quality could be superior in an operation dedicated to one product.

Eventually, the team decided to focus on a relatively non-flexible focused and dedicated large machine. '*You can't imagine the agonies we went through when we decided not to make this a flexible machine. Many of us were not comfortable with saying, "This is going to be a photoresist machine exclusively, and if the market goes away we're in real trouble". We had a lot of debate about that. Eventually, we more or less reached a consensus for focus but it was certainly one of the toughest decisions we ever made.*' (Frank Jansen) The capital cost savings of a focused facility and operating costs savings of up to 25 per cent were powerful arguments, as was the philosophy of total process dedication. '*The key word for us was* ***focus*** .

We wanted to be quite clear about what was needed to satisfy our customer in making this single type of product. As well as providing significant cost savings to us it made it a lot easier to identify the root causes of any problems because we would not have to worry about how it might affect other products. It's all very clear. When the line was down we would not be generating revenue! It would also force us to understand our own performance. At our other operations, if a line goes down, the people can be shifted to other responsibilities. We don't have other responsibilities here – we're either making it or we're not.' (Frank Jansen)

When the Rotterdam operation started producing, the team had tweaked the design to bring the capacity at start-up to 32 MSM per year. And notwithstanding some initial teething troubles it was, from the start, a technical and commercial success. Within 6 months a contract was signed with Phanchem to supply 100 per cent of Phanchem's needs for the next 10 years. Phanchem's decision was based on the combination of manufacturing and business focus that the Rotterdam team has achieved, a point stressed by Frank Jansen. *'Co-locating all necessary departments on the Rotterdam site was seen as particularly important. All the technical functions and the marketing and business functions are now on site.'*

Developing the supply relationship

At the time of the start-up, product produced in Rotterdam was shipped to Phanchem's facility near Frankfurt, Germany, almost 500 km away. This distance caused a number of problems including some damage in transit and delays in delivery. However, the relationship between AE and Phanchem remained sound; helped by the two companies' co-operation during the Rotterdam start-up. *'We had worked closely with them during the design and construction of the new Rotterdam facility. More to the point, they saw that they would certainly achieve cost savings from the plant, with the promise of more savings to come as the plant moved down the learning curve.'* (Frank Jansen) The closeness of the relationship between the two companies was a result of their staff working together. AE engineers were impressed by their customer's willingness to help out while they worked on overcoming the start-up problems. Similarly AE had helped Phanchem when they needed extra supplies at short notice. As Frank Jansen said, *'partly because we worked together on various problems the relationship has grown stronger and stronger.'*

In particular, the idea of a physically closer relationship between AE and Phanchem was explored. *'During the negotiations with Phanchem for our 100 per cent contract there had been some talk about co-location, but I don't think anyone took it particularly seriously. Nevertheless, there was general agreement that it would be a good thing to do. After all, our success as Phanchem's sole supplier of coated photoresist was tied in to their success as a player in the global market; what was good for Phanchem was good for AE.'* (Frank Jansen) Several options were discussed within and between the two companies. Phanchem had, in effect, to choose between four options:

1. Stay where they were near Frankfurt.
2. Relocate to the Netherlands (which would give easier access to port facilities) but not too close to AE (an appropriate site was available 30 km from Rotterdam).
3. Locate to a currently vacant adjacent site across the road from AE's Rotterdam plant.
4. Co-locate within an extension that could be specially built onto the AE plant at Rotterdam.

Evaluating the co-location options

Relatively early in the discussions between the two companies, the option of 'doing nothing' by staying in Frankfurt was discounted. Phanchem wanted to sell their valuable site near Frankfurt. The advantages of some kind of move were significant. The option of Phanchem moving to a site 30 km from Rotterdam was considered but rejected because it had no advantages over locating even closer to the Rotterdam plant. Phanchem also strongly considered building and operating a facility across the road from the Rotterdam plant. But eventually the option of locating in a building attached to AE's Rotterdam operation became the preferred option. Co-location would have a significant impact on Phanchem's competitiveness by reducing their operating costs, enabling them to gain market share by offering quality film at attractive prices, thus increasing volume for AE. The managers at the Rotterdam plant also looked forward to an even closer operational relationship with the customer. *'Initially, there was some resistance in the team to having a customer on the same site as ourselves. No one in AE had ever done it before. The step from imagining our customer across the road to imagining them on the same site took some thinking about. It was a matter of getting used to the idea, taking one step at a time.'* (Frank Jansen)

The customer becomes a paying guest

However, when Frank and the Rotterdam managers presented their proposal for extending the plant to the AE board the proposal was not well received. '*Leasing factory space to our customer seemed a long way from our core business. As one Executive Committee member said, we are manufacturers; we aren't in the real estate business. But we felt that it would be beneficial for both companies*'. (Frank Jansen) And even when the proposal was eventually accepted, there was still concern over sharing a facility. In fact the Executive Committee insisted, that the door between the two companies' areas should be capable of being locked from both sides. Yet the construction and commissioning of the new facility for Phanchem was also a model of cooperation. Now, all visitors to the plant are shown the door that had to be 'capable of being locked from both sides' and asked how many times they think it has been locked. The answer, of course, is 'never'.

QUESTIONS

1. What were the key structure and scope decisions taken by Aarens Electronic?
2. What were the risks involved in adopting a process design that was 'totally dedicated' to the one customer's needs?
3. What were the advantages and disadvantages of each location option open to Phanchem, and why do you think they erventually chose to co-locate with AE?

APPLYING THE PRINCIPLES

Some of these exercises can be answered by reading the chapter. Others will require some general knowledge of business activity and some might require an element of investigation. **Model answers for the first two questions are available on the student companion website.** For model answers to the other questions in this section, please ask your tutor/lecturer.

1. Consider the music business as a supply network. How did music downloads and streaming affect to artists' sales? What implications did electronic music transmission have for record shops?

2. A data centre is 'a facility composed of networked computers and storage that businesses or other organisations use to organise, process, store and disseminate large amounts of data. A business typically relies heavily upon the applications, services and data contained within a data center, making it a focal point and critical asset for everyday operations'.[9] These facilities can contain network equipment, servers, data storage and back-up facilities, software applications for large companies, and more. Very few businesses (or people) do not rely on them. And determining their location is a crucial decision for those operations running them. In fact, such businesses usually have a set method for choosing data centre location.

   Visit the websites of the type of business that run data centres (such as Intel, Cisco or SAP) and devise a set of criteria that could be used to evaluate potential sites.

3. A company that produces concrete paving slabs is introducing a new range of 'textured' non-slip products. To do this, it must invest in a new machine. Demand is forecast to be around 10,000 units per month for the first year and approximately 24,000 units per month after that. The machines that produce these products have a capacity of 10,000 units per month. They have a fixed cost of €20,000 per month and a variable cost of processing of €1 per unit. The company has forecast that they will be able to charge €4 per unit. It has been suggested that they would make higher profits if sales were restricted to 20,000 units per month in the second year. Is this true?

4 The Fast and Efficient (FAC) Transport Group is reviewing its fleet maintenance operations. *'Our lease on our current maintenance and repair facilities site will expire in a year, and we need to decide how to operate in the future. Currently, we have the one site with 5 repair bays. This can cope with our fleet of 40 trucks. But demand is growing; and within 2 or 3 years we hope to be operating around 55–60 trucks. So we will have to choose a site (or sites) that allows for this increase. And that leads me to the next issue – should we stick to operating one central site, or should we plan to have two sites, one for the North and one for the South of our region?'*

As far as FAC's operations managers could forecast, the costs of having one or two sites would be as follows:

One site – Fixed cost of establishing the site = €300,000, variable cost of servicing trucks = €14,000 per truck per year

Two sites – Fixed cost of establishing the sites (for both) = €500,000, variable cost of servicing trucks = €10,000 per truck per year (they will be out of action for less time because sites would be closer).

At what level of demand (in terms of the number of trucks operated by the company) will the two-site proposal be cheaper?

5 'Globalisation is very much a 'mixed blessing'. There is little doubt that it has lifted millions out of poverty, but it has also led to the destruction of traditional cultures in developing countries and many jobs in the developed world.' Draw up lists of what you see as the advantages and disadvantages of globalisation.

Notes on chapter

1 Christopher, M. (2011) 'Logistics and supply chain management', *Financial Times,* London.

2 Sources include: Burton, G. (2013) 'ARM vs Intel: a battle of business models', *Computing* 29 May; Turley, J. (2014) 'Intel vs. ARM: Two titans' tangled fate', InfoWorld.com, Created 27 February.

3 Sources include: Butler, S. (2014) 'Rana Plaza disaster marked by Oxford Street demonstration', *The Guardian,* Thursday 24 April; *Economist* (2013) 'Disaster at Rana Plaza - A gruesome accident should make all bosses think harder about what behaving responsibly means', 4 May; *Economist* (2013) 'Bangladesh's clothing industry- Bursting at the seams', 26 October.

4 For further details of the origin of this concept see: Moore, J. F. (1993) 'Predators and prey: a new ecology of competition', *Harvard Business Review,* May–June.

5 Moore, J. F. (1993) op. cit.

6 Bacon, G. Machan, I. and Dnyse, J. (2008) 'Offshore challenges, manufacturing', *Institute of Electrical Engineers,* January.

7 Sources include: Blakely, R. (2010)' Britain can learn from India's assembly-line heart operations, says doctor', *The Times,* 14 May; *Economist* (2011) 'Economies of scale made steel - The economics of very big ships', 12 November.

8 Sources include: Arlidge, J. (2013) 'Rolls finds its Derby in the east', *Sunday Times,* 27 October; Raghuvanshi, G. (2013) 'Rolls-Royce Pushes Focus on Singapore', *Wall Street Journal,* 15 September; Syed, S. (2013) 'Rolls-Royce gears up for Singapore production', *BBC News* website 21 February.

9 Definition from techtarget.com, see searchdatacenter.techtarget.com/

TAKING IT FURTHER

Carmel, E. and Tija, P. (2005) *Offshoring Information Technology: Sourcing and Outsourcing to a Global Workforce*, ***Cambridge University Press***. *An academic book on outsourcing.*

Chopra, S. and Meindl, P. (2009) *Supply Chain Management: Strategy, planning and operations* (4th edn), ***Prentice Hall***. *A good textbook that covers both strategic and operations issues.*

Corbett, M. F. (2010) *The Outsourcing Revolution: Why it Makes Sense and How to Do it Right*, ***Kaplan***. *Not an academic book on outsourcing.*

Gunasekarana, A., Kee-hung Laib, K. and Edwin Cheng, T.C. (2008) 'Responsive supply chain: A competitive strategy in a networked economy', *Omega*, ***vol. 36, Issue 4***.

Moore, J. F. (2013) *Shared Purpose: A thousand business ecosystems, a connected community, and the future*, ***CreateSpace Independent Publishing Platform***. *A short but sweet progress report on a study of ARM Holdings (see example at the beginning of the chapter) and its 1,000-plus community of partners from the founder of the business ecosystems idea.*

Oshri, I., Kotlarsky J. and Willcocks, L. P. (2015) *The Handbook of Global Outsourcing and Offshoring* (3rd edn a), ***Palgrave Macmillan***. *All you need to know.*

Singh Srai, J. and Gregory, M. (2008) 'A supply network configuration perspective on international supply chain development', *International Journal of Operations & Production Management*, vol. 28, Issue 5.

Steger, M. B. (2017) *Globalization: A Very Short Introduction*, ***OUP Oxford***. *Not just a business book, also covers political, cultural, ideological, and environmental issues.*

Vashistha, A. and Vashistha, A. (2006) *The Offshore Nation: Strategies for Success in Global Outsourcing and Offshoring*, ***McGraw Hill Higher Education***. *Another topical book on outsourcing*

SUPPLEMENT TO CHAPTER 4

Forecasting

Introduction

Some forecasts are accurate. We know exactly what time the sun will rise at any given place on Earth tomorrow, or one day next month, or even next year. Forecasting in a business context, however, is much more difficult and therefore prone to error. We do not know precisely how many orders we will receive or how many customers will walk through the door tomorrow, next month, or next year. Such forecasts, however, are necessary to help managers make decisions about resourcing the organisation for the future.

Forecasting – knowing the options

Simply knowing that demand for your goods or services is rising or falling is not enough in itself. Knowing the rate of change is likely to be vital to business planning. A firm of lawyers may have to decide the point at which, in their growing business, they will have to take on another partner. Hiring a new partner could take months so they need to be able to forecast when they expect to reach that point and then when they need to start their recruitment drive. The same applies to a plant manager who will need to purchase new plant to deal with rising demand. She may not want to commit to buying an expensive piece of machinery until absolutely necessary but in enough time to order the machine and have it built, delivered, installed and tested. The same is so for governments, whether planning new airports or runway capacity, or deciding where and how many primary schools to build.

The first question is to know how far you need to look ahead and this will depend on the options and decisions available to you. Take the example of a local government, where the number of primary age children (5–11 years) is increasing in some areas and declining in other areas within its boundaries. It is legally obliged to provide school places for all such children. Government officials will have a number of options open to them and they may each have different lead times associated with them. One key step in forecasting is to know the possible options and the lead times required to bring them about (see Table 4.3)

Table 4.3 Options available and lead time required for dealing with changes in numbers of school children

Options available	*Lead time required*
Hire short-term teachers	Hours
Hire staff	↓
Build temporary classrooms	
Amend school catchment areas	
Build new classrooms	
Build new schools	Years

Individual schools can hire (or lay-off) short term (supply) teachers from a pool, not only to cover for absent teachers, but also to provide short-term capacity while teachers are hired to deal with increases in demand. Acquiring (or dismissing) such temporary cover may only require a few hours' notice. (This is often referred to as short-term capacity management.)

Hiring new or laying-off existing staff is another option, but both of these may take months to complete. (Medium-term capacity management.)

A shortage of accommodation may be fixed in the short to medium term by hiring or buying temporary classrooms. It may only take a couple of weeks to hire such a building and equip it ready for use.

It may be possible to amend catchment areas between schools to try to balance an increasing population in one area against a declining population in another. Such changes may require lengthy consultation processes.

In the longer term, new classrooms or even new schools may have to be built. The planning, consultation, approval, commissioning, tendering, building and equipping process may take 1 to 5 years depending on the scale of the new build.

Knowing the range of options, managers can then decide the time scale for their forecasts; indeed several forecasts might be needed for the short term, medium term and long term.

In essence forecasting is simple

In essence forecasting is easy. To know how many children may turn up in a local school tomorrow, you can use the number that turned up today. In the long term, in order to forecast how many primary-aged children will turn up at a school in 5 years' time, one need simply look at the birth statistics for the current year for the school's catchment area. See Figure 4.10.

However, such simple extrapolation techniques are prone to error and indeed such approaches have resulted in some local governments committing themselves to building schools which five or six years later, when complete, had few children, while other schools were bursting at the seams with temporary classrooms and temporary teachers, often resulting in falling morale and declining educational standards. The contextual variables (see Figure 4.11) that will have a potentially significant impact on, for example, the school population five years hence. For example:

One minor factor in developed countries, though a major factor in developing countries, might be the death rate in children between birth and 5 years of age. This may be dependent upon location with a slightly higher mortality rate in the poorer areas, compared to the more affluent areas.

Another more significant factor is immigration and emigration, as people move into or out of the local area. This will be affected by housing stock and housing developments and the ebb and flow of jobs in the area and the changing economic prosperity in the area.

One key factor that has an impact on the birth rate in an area is the amount and type of the housing stock. City centre tenement buildings tend to have a higher proportion of children per dwelling, for example, than suburban semi-detached houses. So not only will existing housing

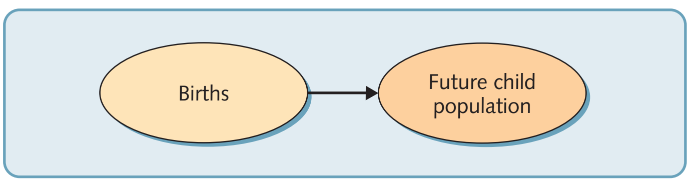

Figure 4.10 Simple prediction of future child population

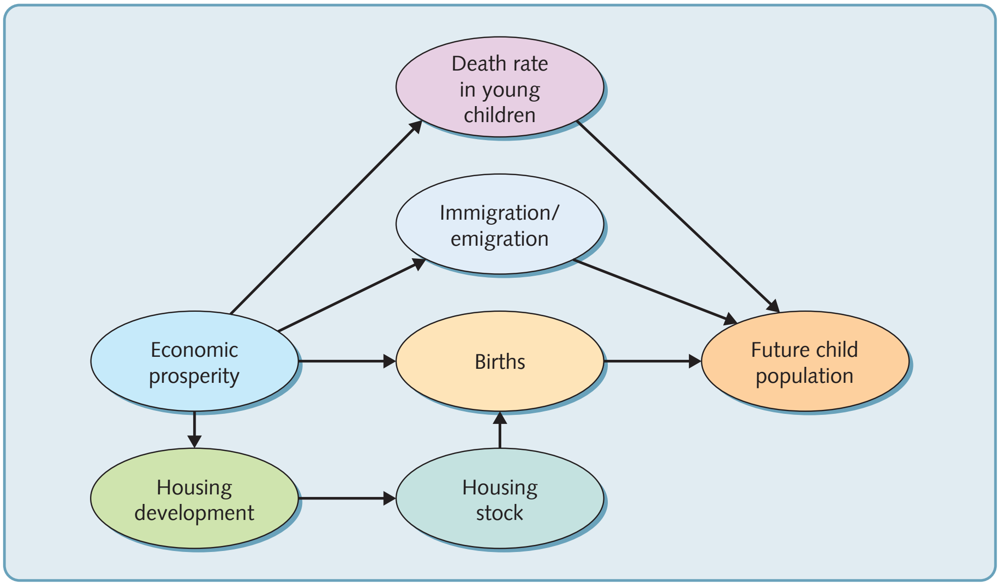

Figure 4.11 **Some of the key causal variables in predicting child populations**

stock have an impact on the child population, but so also will the type of housing developments under construction, planned and proposed.

Approaches to forecasting

There are two main approaches to forecasting. Managers sometimes use qualitative methods based on opinions, past experience and even best guesses. There is also a range of qualitative forecasting techniques available to help managers evaluate trends, causal relationships and make predictions about the future. In addition, quantitative forecasting techniques can be used to model data. Although no approach or technique will result in an accurate forecast, a combination of qualitative and quantitative approaches can be used to great effect by bringing together expert judgements and predictive models.

Qualitative methods

Imagine you were asked to forecast the outcome of a forthcoming football match. Simply looking at the teams' performance over the last few weeks and extrapolating it is unlikely to yield the right result. Like many business decisions the outcome will depend on many other factors. In this case the strength of the opposition, their recent form, injuries to players on both sides, the match location and even the weather will have an influence on the outcome. A qualitative approach involves collecting and appraising judgements, options, even best guesses as well as past performance from 'experts' to make a prediction. There are several ways this can be done: a panel approach, Delphi method and scenario planning.

Panel approach

Just as panels of football pundits gather to speculate about likely outcomes, so too do politicians, business leaders, stock market analysts, banks and airlines. The panel acts like a focus group, allowing everyone to talk openly and freely. Although there is the great advantage of several brains being better than one, it can be difficult to reach a consensus, or sometimes the views of the loudest or highest status may emerge (the bandwagon effect). Although more reliable than one person's views the panel approach still has the weakness that everybody, even the experts, can get it wrong.

Delphi method

Perhaps the best-known approach to generating forecasts using experts is the Delphi method.[1] This is a more formal method[1] which attempts to reduce the influences from the procedures of face-to-face meetings. It employs a questionnaire, emailed or posted to the experts. The replies are analysed and summarised and returned, anonymously, to all the experts. The experts are then asked to reconsider their original response in the light of the replies and arguments put forward by the other experts. This process is repeated several more times to conclude either with a consensus, or at least a narrower range of decisions. One refinement of this approach is to allocate weights to the individuals and their suggestions based on, for example, their experience, their past success in forecasting, other people's views of their abilities. The obvious problems associated with this method include constructing an appropriate questionnaire, selecting an appropriate panel of experts and trying to deal with their inherent biases.

Scenario planning

One method for dealing with situations of even greater uncertainty is scenario planning. This is usually applied to long-range forecasting, again using a panel. The panel members are usually asked to devise a range of future scenarios. Each scenario can then be discussed and the inherent risks considered. Unlike the Delphi method, scenario planning is not necessarily concerned with arriving at a consensus; it looks at the possible range of options and putting plans in place to try to avoid the ones that are least desired and taking action to follow the most desired.

Quantitative methods

There are two main approaches to quantitative forecasting: time series analysis; and causal modelling techniques:

1. Time series examine the pattern of past behaviour of a single phenomenon over time, taking into account reasons for variation in the trend in order to use the analysis to forecast the phenomenon's future behaviour.
2. Causal modelling is an approach which describes and evaluates the complex cause–effect relationships between the key variables (such as in Figure 4.11).

Time series analysis

Simple time series plot a variable over time then by removing underlying variations with assignable causes use extrapolation techniques to predict future behaviour. The key weakness with this approach is that it simply looks at past behaviour to predict the future, ignoring causal variables which are taken into account in other methods such as causal modelling or qualitative techniques. For example, suppose a company is attempting to predict the future sales of a product. The past three years' sales, quarter by quarter, are shown in Figure 4.12(a). This series of past sales may be analysed to indicate future sales. For instance, underlying the series might be a linear upward trend in sales. If this is taken out of the data, as in Figure 4.12(b), we are left with a cyclical seasonal variation. The mean deviation of each quarter from the trend line

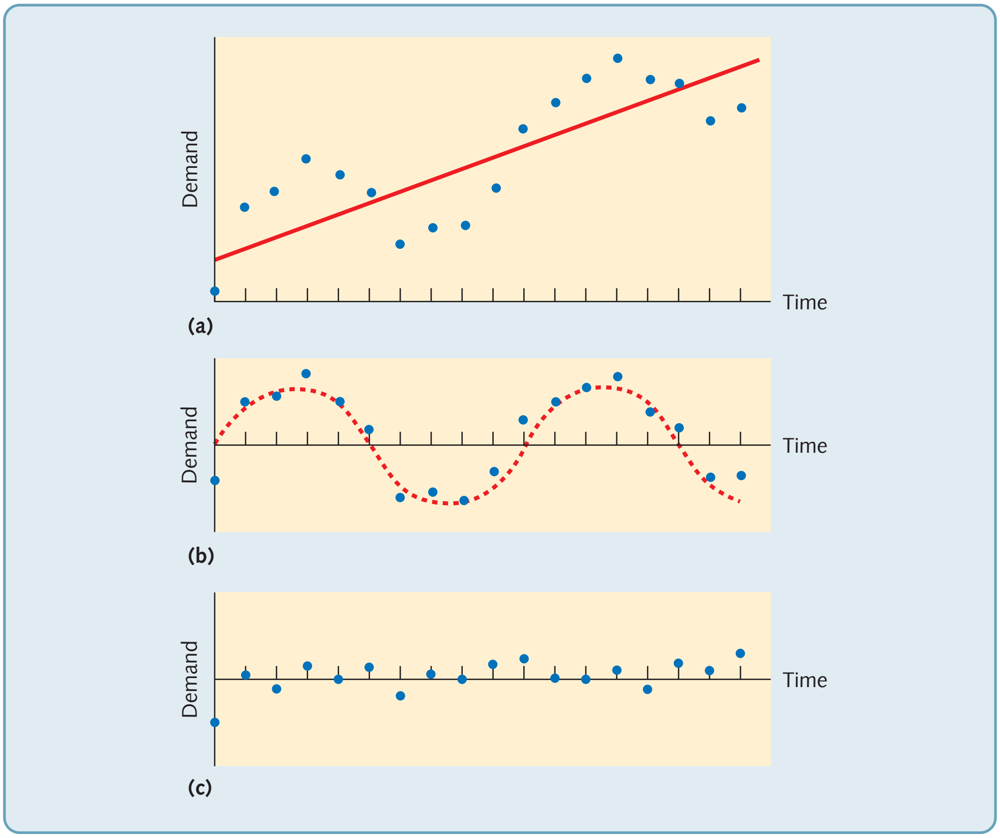

Figure 4.12 Time series analysis with (a) trend, (b) seasonality and (c) random variation

can now be taken out, to give the average seasonality deviation. What remains is the random variation about the trends and seasonality lines, Figure 4.12(c). Future sales may now be predicted as lying within a band about a projection of the trend, plus the seasonality. The width of the band will be a function of the degree of random variation.

Forecasting un-assignable variations

The random variations which remain after taking out trend and seasonal effects are without any known or assignable cause. This does not mean that they do not have a cause, however, just that we do not know what it is. Nevertheless, some attempt can be made to forecast it, if only on the basis that future events will, in some way, be based on past events. We will examine two of the more common approaches to forecasting which are based on projecting forward from past behaviour. These are:

- moving-average forecasting
- exponentially smoothed forecasting.

The moving-average approach to forecasting takes the previous *n* periods' actual demand figures, calculates the average demand over the *n* periods and uses this average as a forecast for the next period's demand. Any data older than the *n* periods play no part in the next period's forecast. The value of *n* can be set at any level, but is usually in the range 4 to 7.

EXAMPLE

Eurospeed parcels

Table 4.4 shows the weekly demand for Eurospeed, a European-wide parcel delivery company. It measures demand, on a weekly basis, in terms of the number of parcels which it is given to deliver (irrespective of the size of each parcel). Each week, the next week's demand is forecast by taking the moving average of the previous four weeks' actual demand. Thus if the forecast demand for week *t* is *Ft* and the actual demand for week *t* is *At,* then:

$$F_t = \frac{1}{4}(A_{t-4} + A_{t-3} + A_{t-2} + A_{t-1})$$

For example, the forecast for week 35:

$$F_{35} = (72.5 + 66.7 + 68.3 + 67.0)/4$$
$$= 68.8$$

Exponential smoothing

There are two significant drawbacks to the moving-average approach to forecasting. First, in its basic form, it gives equal weight to all the previous *n* periods which are used in the calculations (although this can be overcome by assigning different weights to each of the *n* periods). Second, and more important, it does not use data from beyond the *n* periods over which the moving average is calculated. Both these problems are overcome by exponential smoothing, which is also somewhat easier to calculate. The exponential-smoothing approach forecasts demand in the next period by taking into account the actual demand in the current period and

Table 4.4 Moving-average forecast calculated over a four-week period

Week	*Actual demand (thousands)*	*Forecast*
20	63.3	
21	62.5	
22	67.8	
23	66.0	
24	67.2	64.9
25	69.9	65.9
26	65.6	67.7
27	71.1	66.3
28	68.8	67.3
29	68.4	68.9
30	70.3	68.5
31	72.5	69.7
32	66.7	70.0
33	68.3	69.5
34	67.0	69.5
35		68.6

the forecast which was previously made for the current period. It does so according to the formula:

$$F_t = \alpha A_{t-1} + (1 - x)F_{t-1}$$

where α = the smoothing constant.

The smoothing constant α is, in effect, the weight which is given to the last (and therefore assumed to be most important) piece of information available to the forecaster. However, the other expression in the formula includes the forecast for the current period which included the previous period's actual demand, and so on. In this way all previous data have a (diminishing) effect on the next forecast.

Table 4.5 shows the data for Eurospeed's parcels forecasts using this exponential-smoothing method, where $\alpha = 0.2$. For example, the forecast for week 35 is:

$$F_{35} = (0.2 \times 67.0) + (0.8 \times 68.3) = 68.04$$

The value of α governs the balance between the responsiveness of the forecasts to changes in demand, and the stability of the forecasts. The closer α is to 0 the more forecasts will be dampened by previous forecasts (not very sensitive but stable). Figure 4.13 shows the Eurospeed volume data plotted for a four-week moving average, exponential smoothing with $\alpha = 0.2$ and exponential smoothing with $\alpha = 0.3$.

Table 4.5 Exponentially smoothed forecast calculated with smoothing constant $\alpha = 0.2$

Week (t)	*Actual demand (thousands) (A)*	*Forecast (F1 = can't do this)*
20	63.3	60.00
21	62.5	60.66
22	67.8	60.03
23	66.0	61.58
24	67.2	62.83
25	69.9	63.70
26	65.6	64.94
27	71.1	65.07
28	68.8	66.28
29	68.4	66.78
30	70.3	67.12
31	72.5	67.75
32	66.7	68.70
33	68.3	68.30
34	67.0	68.30
35		68.04

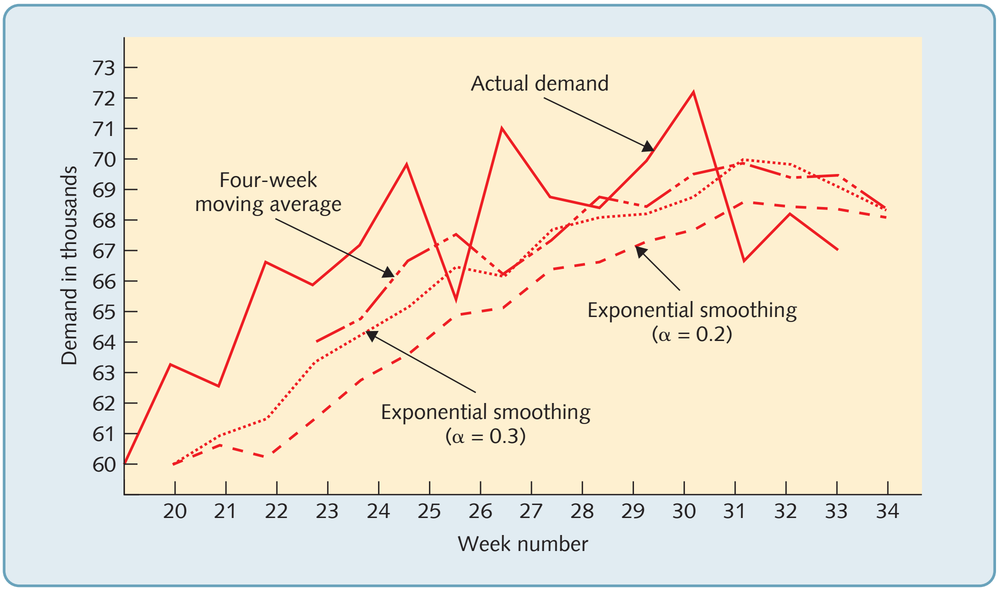

Figure 4.13 A comparison of a moving-average forecast and exponential smoothing with the smoothing constant $\alpha = 0.2$ and 0.3

Causal models

Causal models often employ complex techniques to understand the strength of relationships between the network of variables and the impact they have on each other. Simple regression models try to determine the 'best fit' expression between two variables. For example, suppose an ice-cream company is trying to forecast its future sales. After examining previous demand, it figures that the main influence on demand at the factory is the average temperature of the previous week. To understand this relationship, the company plots demand against the previous week's temperatures. This is shown in Figure 4.14. Using this graph, the company can make a reasonable prediction of demand, once the average temperature is known, provided that the other conditions prevailing in the market are reasonably stable. If they are not, then these other factors which have an influence on demand will need to be included in the regression model, which becomes increasingly complex.

These more complex networks comprise many variables and relationships each with their own set of assumptions and limitations. While developing such models and assessing the importance of each of the factors and understanding the network of interrelationships is beyond the scope of this text, many techniques are available to help managers undertake this more complex modelling and also feedback data into the model to further refine and develop it, in particular structural equation modelling.

The performance of forecasting models

Forecasting models are widely used in management decision-making, and indeed most decisions require a forecast of some kind, yet the performance of this type of model is far from impressive. Hogarth and Makridakis,[2] in a comprehensive review of the applied management and finance literature, show that the record of forecasters using both judgement and sophisticated mathematical methods is not good. What they do suggest, however, is that certain forecasting techniques perform better under certain circumstances. In short-term forecasting there is:

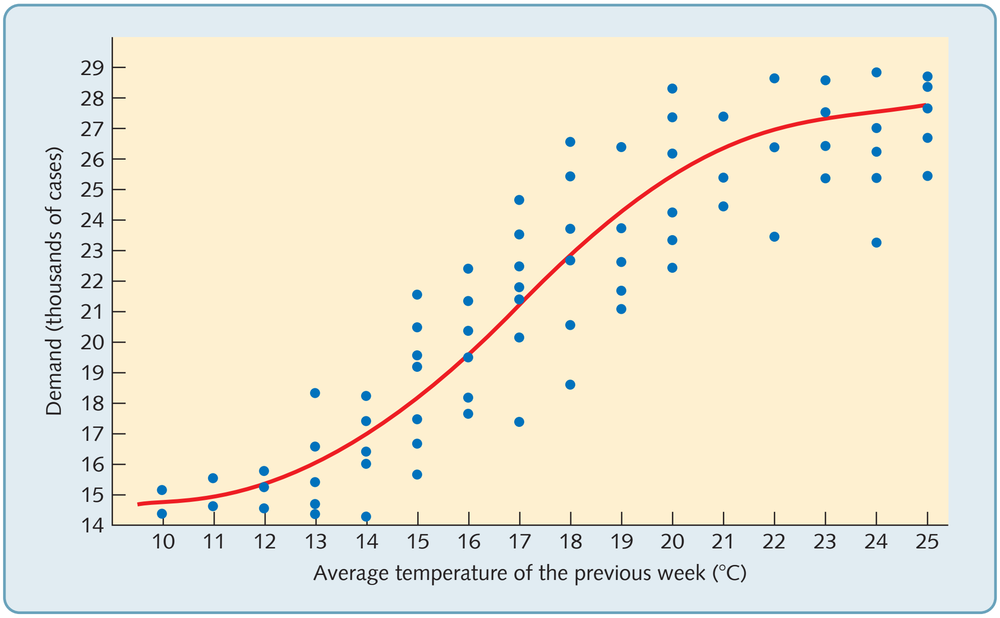

Figure 4.14 **Regression line showing the relationship between the previous week's average temperature and demand**

. . . considerable inertia in most economic and natural phenomena. Thus the present states of any variables are predictive of the short-term future (i.e. 3 months or less). Rather simple mechanistic methods, such as those used in time series forecasts, can often make accurate short-term forecasts and even out-perform more theoretically elegant and elaborate approaches used in econometric forecasting.[3]

Long-term forecasting methods, although difficult to judge because of the time lapse between the forecast and the event, do seem to be more amenable to an objective causal approach. In a comparative study of long-term market forecasting methods, Armstrong and Grohman[4] conclude that econometric methods offer more accurate long-range forecasts than do expert opinion or time series analysis, and that the superiority of objective causal methods improves as the time horizon increases.

Notes on chapter supplement

1 Linstone, H.A. and Turoof, M. (1975) *The Delphi Method: Techniques and Applications*, Addison-Wesley.

2 Hogarth, R.M. and Makridakis, S. (1981) 'Forecasting and planning: an evaluation', *Management Science*, vol. 27, 115–38.

3 Hogarth, R.M. and Makridakis, S., op. cit.

4 Armstrong, J.S. and Grohman, M.C. (1972) 'A comparative study of methods for long-range market forecasting', *Management Science,* vol. 19, no. 2, 211–21.

TAKING IT FURTHER

Hyndman, R. J. and Athanasopoulos, G. (2013) ***Forecasting: principles and practice, OTexts***. *A textbook that provides a comprehensive introduction to forecasting methods.*

Morlidge, S. and Player, S (2010) ***Future Ready: How to Master Business Forecasting*****, John Wiley & Sons.** *A good, not too technical, treatment.*

Silver, N. (2013) ***The Signal and the Noise: The Art and Science of Prediction, Penguin.*** *Nate Silver is probably the best-known forecaster in the world (he accurately predicted the results of every single state in the 2012 US election), and for good reason. A very readable book.*

Tetlock, P. and Gardner, D (2016) Superforecasting: ***The Art and Science of Prediction, Random House Books.*** *Readable and sensible.*

5

Process design 1 – Positioning

Introduction

In Chapter 1 we described a 'process' as an arrangement of resources and activities that transform inputs into outputs that satisfy (internal or external) customer needs. And processes are everywhere. They are the building blocks of all operations, and their design will affect the performance of the whole operation and, eventually, the contribution it makes to its supply network. No one, in any function or part of the business, can fully contribute to its competitiveness if the processes in which they work are poorly designed and ineffective. It is not surprising then that process design has become such a popular topic in the management press and among consultants. This chapter is the first of two that examine the design of processes. To understand the difference between this chapter (positioning) and the following one (analysis), go back to our definition of a process, 'an arrangement of *resources* and *activities*. . .'. This chapter is primarily concerned with the resources in processes and, more specifically, how process resources must reflect the volume and variety requirements placed on them. The next chapter examines the activities within processes, and how they can be analysed in order to understand better how they will operate and, therefore, how their performance could be improved. Figure 5.1 shows the position of the ideas described in this chapter in the general model of operations management.

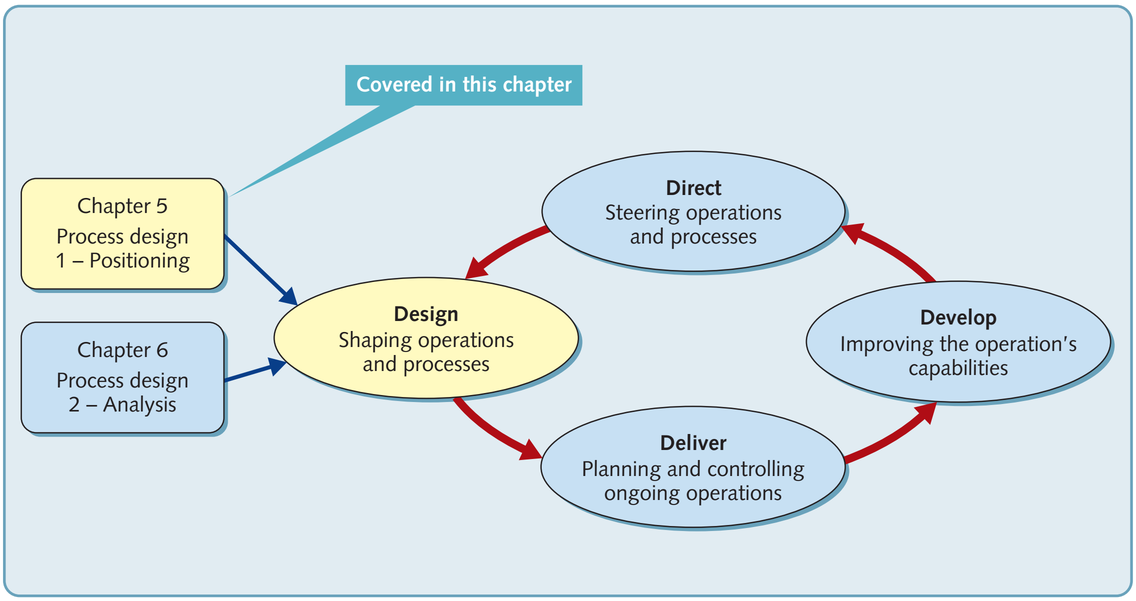

Figure 5.1 **Process design – positioning is concerned with conceiving the nature of the resources that make up the process so that they are appropriate for their volume–variety position**

EXECUTIVE SUMMARY

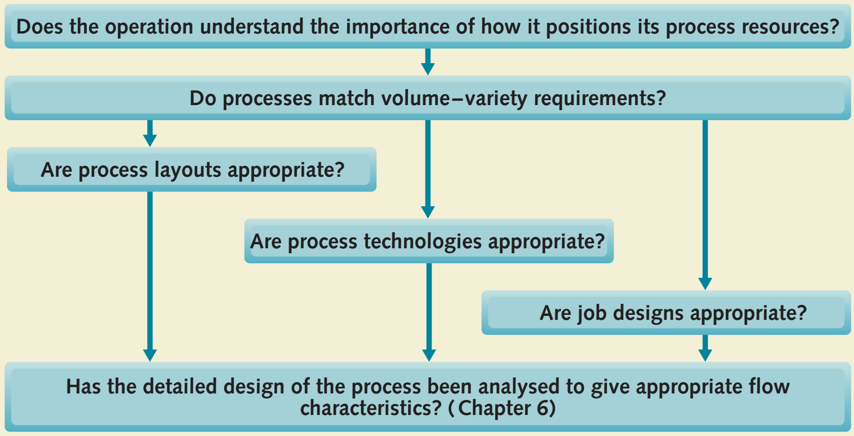

Does the operation understand the importance of how it positions its process resources?

Process design is concerned with conceiving the nature of the resources that make up the process and their detailed workings. The first task of process design is to conceive the nature of the resources that make up the process and how they are arranged. Without appropriate resources it is difficult (maybe impossible) for the process ever to operate as effectively as it could do. A useful way to do this is to position the process in terms of its volume and variety characteristics. Only after this stage, should the second task (of conceiving the detailed workings of the process) be attempted.

Do processes match volume–variety requirements?

Volume and variety are particularly influential in the design of processes. They also tend to go together in an inverse relationship. High-variety processes are normally low volume and vice versa. So processes can be positioned on the spectrum between low volume and high variety and high volume and low variety. At different points on this spectrum, processes can be described as distinct process 'types'. Different terms are used in manufacturing and service to identify these types. Working from low volume and high variety towards high volume and low variety, the process types are: project processes; jobbing processes; batch processes, mass processes; and continuous processes. Working in the same sequence the service types are known as: professional services; service shops; and mass services. Whatever terminology is used, the overall design of the process must fit its volume–variety position. This is usually summarised in the form of the 'product–process' matrix.

Are process layouts appropriate?

There are different ways in which the different resources within a process (people and technology) can be arranged relative to each other. But however this is done it should reflect the process's volume–variety position. Again, there are pure 'types' of layout that correspond with the different volume–variety positions. These are: fixed-position layout; functional layout; cell layout; and product layout. Many layouts are hybrids of these pure types, but the type chosen is influenced by the volume and variety characteristics of the process.

Are process technologies appropriate?

Process technologies are the machine's equipment and devices that help processes to transform materials, information and customers. It is different to product technology that is embedded within the product or service itself. Again, product technology should reflect volume and variety. In particular, the degree of automation in the technology, the scale and/or scalability of the technology and the coupling and/or connectivity of the technology should be appropriate to volume and variety. Generally, low volume and high variety requires relatively unautomated, general-purpose, small-scale and flexible technologies. By contrast, high-volume and low-variety processes require automated, dedicated and large-scale technologies that are sometimes relatively inflexible.

Are job designs appropriate?

Job design is about how people carry out their tasks within a process. It is particularly important because it governs people's expectations and perceptions of their contribution to the organisation, as well as being a major factor in shaping the culture of the organisation. Some aspects of job design are common to all processes, irrespective of their volume and variety position. These are such things as ensuring the safety of everyone affected by the process, ensuring a firm ethical stance and upholding an appropriate work/life balance. However, other aspects of job design are influenced by volume and variety. In particular, the extent of division of labour, the degree to which jobs are defined and the way in which job commitment is encouraged. Broadly, high-variety and low-volume processes require broad, relatively undefined jobs with decision-making discretion. Such jobs tend to have intrinsic job commitment. By contrast, high-volume and low-variety processes tend to require jobs that are relatively narrow in scope and closely defined with relatively little decision-making discretion. This means some deliberative action is needed in the design of the job (such as job enrichment) in order to help maintain commitment to the job.

DIAGNOSTIC QUESTION

Does the operation understand the importance of how it positions its process resources?

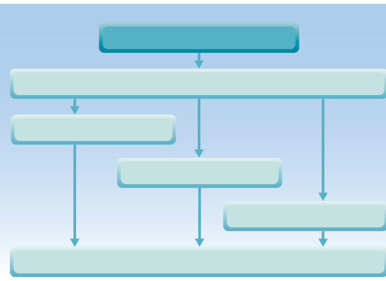

'Positioning' processes is the first step in making sure that they are well designed. To 'design' is to conceive the looks, arrangement and workings of something *before it is constructed.* In that sense design is a conceptual exercise. Yet it is also one that must deliver a solution that will work in practice. Design is also an activity that can be approached at different levels of detail. One may envisage the general shape and intention of something before getting down to defining its details. However, it is often only through getting to grips with the detail of a design that the feasibility of its overall shape can be assessed. So it is with designing processes. First, one must consider the overall shape and nature of the resources that make up the process. Second, one must analyse the details of the activities within the process in order to ensure that it fulfils its objectives effectively. But don't think of this as a simple sequential process. There may be aspects concerned with the broad positioning of the process that will need to be modified following its more detailed analysis.

The important point here is that if one gets the 'high-level' design wrong, it is difficult to finish up with a well-designed process. Processes with the wrong resources; unsuitable technology, staffed by people with inappropriately designed jobs, or resources arranged in an unsuitable manner, cannot ever realise their full potential. And the key to deciding how to position process resources is the volume and variety characteristics under which it has to operate. Which is why, in this chapter, we discuss the more general approach to process design by showing how a process's position on the volume–variety scale will influence its layout, technology and the design of its jobs. In the next chapter we will discuss the more detailed aspects of process design, in particular its objectives, current configuration, capacity and variability. This is illustrated in Figure 5.2.

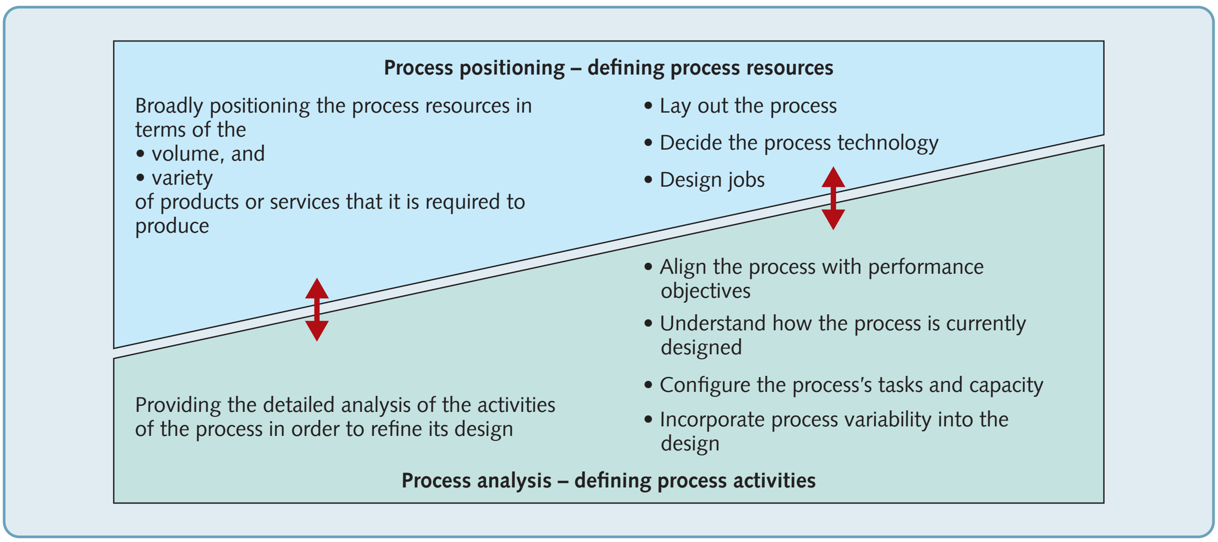

Figure 5.2 Process design is treated in two parts, Positioning, that sets the broad characteristics of the resources in the process, and Analysis, that refines the details of the design by examining the activities within the process

EXAMPLE

Supermarket layout[1]

Successful supermarkets know that the design and layout of their stores has a huge impact on their profitability. Which is why all big supermarket chains do extensive research into the most effective ways of laying out each part of their supermarkets and how best to position specific items. They must maximise their revenue per square metre and minimise the costs of operating the store, while keeping customers happy. At a basic level, supermarkets have to get the amount of space allocated to the different areas right and provide appropriate display and storage facilities. But it is not just the needs of the products that are being sold that are important; the supermarket must also take into consideration the psychological effect that the layout has on their customers. The circulation of customers through the store must be right, but this can vary depending on which country you live in. In the USA shoppers like to work their way through a supermarket anti-clockwise. So, putting the entrance on the right and the checkout on the left encourages this circulation (it is claimed that customers in the US typically spend $2 more per trip when circulating anti-clockwise). By contrast, shoppers in the UK like to shop clockwise, with entrances on the left and checkouts on the right. Why this is so is a bit of a mystery. Some put it down to different driving patterns of each country.

Aisles are made wide to ensure a relatively slow flow of trolleys so that customers pay more attention to the products on display (and buy more). However, wide aisles can come at the expense of reduced shelf space that would allow a wider range of products to be stocked. The actual location of all the products is a critical decision, directly affecting the convenience to customers, their level of spontaneous purchase and the cost of filling the shelves. Although the majority of supermarket sales are packaged tinned or frozen goods, the displays of fruit and vegetables are usually located adjacent to the main entrance, as a signal of freshness and wholesomeness, providing an attractive and welcoming point of entry. Basic products that figure on most people's shopping lists, such as flour, meat sugar and bread, may be spread out towards the back of the store and apart from each other so that customers have to walk along more aisles, passing higher margin items as they search. High-margin items are usually put at eye level on shelves (where they are more likely to be seen) and low-margin products lower down or higher up. Some customers also go a few paces up an aisle before they start looking for what they need. Some supermarkets call the shelves occupying the first metre of an aisle 'dead space', not a place to put impulse-bought goods. But the prime site in a supermarket is the 'gondola-end', the shelves at the end of the aisle. Moving products to this location can increase sales by 200 or 300 per cent. It is not surprising that suppliers are willing to pay for their products to be located here.

However, shopping is changing and supermarket layouts are also changing. Most people still do their food shopping in supermarkets, but are increasingly doing smaller shops, more often. Which is why one large UK supermarket, Sainsbury's, in addition to having a successful and growing convenience store estate, is also testing how it can cater to what it calls specific shopping 'missions' in its supermarkets. This includes a new 'Food to Go' section at the front of the store next to the checkouts. The retailer has also moved its fresh bakery products to this section making it more convenient for people wanting to buy fresh products quickly. Mike Coupe, Chief Executive Officer, Sainsbury's, said: 'This trial is about seeing how far we can go in catering for every shopping mission, whether someone wants to pop in quickly to buy a sandwich for lunch, or whether they have more time and want inspiration for the home, or advice on tech and gadgets. No matter what customers are buying, we know that everyone wants to check out as quickly as possible and giving customers more checkout options to suit them is key to the trial'.

EXAMPLE

Space4 housing processes[2]

You don't usually build a house this way. It's more like the way you would expect an automobile to be made. Nevertheless, Space4's huge building in Birmingham (UK) contains what could be the future of house building. It is a production line whose 90 operators, many of whom have

car assembly experience, are capable of producing the timber-framed panels that form the shell of the new homes at a rate of a house every hour. The automated, systems within the production process control all facets of the operation, ensuring that scheduling and operations are timely and accurate. There is a direct link between the Computer Aided Design (CAD) systems that design the houses and the manufacturing processes that make them, reducing the time between design and manufacture. The machinery itself incorporates automatic predictive and preventative maintenance routines that minimise the chances of unexpected breakdowns. But not everything about the process relies on automation. Because of their previous car assembly experience, staff are used to the just-in-time high-efficiency culture of modern mass production. After production, the completed panels are stacked in 3-metre-high piles and are then fork-lifted into trucks where they are dispatched to building sites across the UK. Once the panels arrive at the building site, the construction workforce can assemble the exterior of a new home in a single day. Because the external structure of a house can be built in a few hours, and enclosed in a weatherproof covering, staff working on the internal fittings of the house, such as plumbers and electricians can have a secure and dry environment in which to work, irrespective of external conditions. Furthermore, the automated production process uses a type of high-precision technology that means there are fewer mistakes in the construction process on site. This means that the approval process from the local regulatory authority takes less time. This process, says Space4, speeds up the total building time from 12–14 weeks to 8–10 weeks.

Also driving the adoption of mass processing were the tough new energy targets. Managing director of Space4, Craig Hagan said: ‘To make the new, improved, standards as part of the drive towards zero-carbon homes need to be energy efficient and insulated. Our success has been down to a combination of the new energy standards favouring our product and the fact we work so quickly.’

Space4 also claims many other advantages for its process. It combines accurate quality conformance, flexibility, a construction routine that uses considerably less water than conventional building methods, which significantly reduces the drying out period and surface cracking, and reduced site-generated waste, requiring fewer skips and providing a tidier, safer site. When built, the homes also have good thermal insulation. The panels from the factory are injected with a special resin mixture that creates a foam that keeps heat in and energy bills down. In fact, sustainability gains are important to the future of this type of process. A report by the Ellen MacArthur Foundation highlighted several potential environmental benefits from off-site construction, including more energy efficient homes. In addition, off-site mass production is also a process that can produce at high volume, while allowing the company to respond quickly to volatile economic conditions because the process has the ability to change volumes relatively easily. Moreover, the quality of the homes can be at least as good as traditionally built homes. ‘They look exactly like a traditional home,’ said Karl Whiteman, executive director of another building firm’s (Berkeley) modular division.

‘The customer will not be able to tell the difference. Our ambition is that the finished quality will be as good as traditionally built homes that have been built on site. I liken it to the evolution of car manufacture: at one time all cars were being hand-built, but now cars are designed for the customer, then designed for manufacture and then assembled in a factory.’

What do they have in common?

Both of the operations described in the two examples illustrate the importance of how an operation decides how to organise and position its productive resources. Supermarkets know that relatively small changes in the layout of their stores can profoundly affect their revenue, while builders like Space4 see significant advantages in rethinking the whole basis of their production processes. Not only that, but both examples show how processes need to be changed as the markets that they serve change. Sainsbury's is responding to changes in its customers' 'shopping missions', Space4 sees the need for a higher volume of housing as well as an increasing acceptance of innovative building methods. Also, both examples have something to tell us about the influence of the volume and variety of demand on the way processes are designed. The volume and variety of demand that a process has to serve (or how it is interpreted by the operation) has a profound effect on the design of the processes. Just think about the design of the processes that would be used in a small 'corner' shop, or for a small handyman/builder. Because they have different volume/variety characteristics, their processes would be designed in very different ways.

DIAGNOSTIC QUESTION

Do processes match volume–variety requirements?

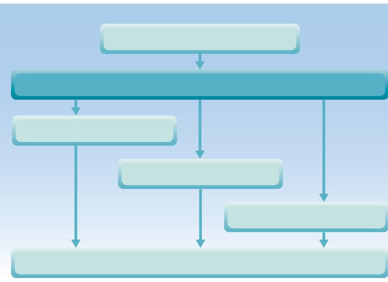

Two factors are particularly important in process design; these are the *volume* and *variety* of the products and services that it processes. Moreover volume and variety are related in so much as low-volume operations processes often have a high variety of products and services, and high-volume operations processes often have a narrow variety of products and services. So, we can *position* processes on a continuum from those that operate under conditions of low volume and high variety, through to those that operate under conditions of high volume and low variety. The volume–variety position of a process influences almost every aspect of its design. Processes with different volume–variety positions will be arranged in different ways, have different flow characteristics and have different technology and jobs. So, a first step in process design is to understand how volume and variety shape process characteristics, and to check whether processes have been configured in a manner that is appropriate for their volume–variety position.

OPERATIONS PRINCIPLE

The design of any process should be governed by the volume and variety it is required to produce.

The 'product–process' matrix

The most common method of illustrating the relationship between a process's volume–variety position and its design characteristics is shown in Figure 5.3. Often called the 'product–process' matrix, it can in fact be used for any type of process whether producing products or services. The underlying idea of the product–process matrix is that many of the more important elements of process design are strongly related to the volume–variety position of the process. So, for any process, the tasks that it undertakes, the flow of items through the process, the layout of its resources, the technology it uses and the design of jobs, are all strongly influenced by its volume–variety position. This means that most processes should lie close to the diagonal of the matrix that represents the 'fit' between the process and its volume–variety position. This is called the 'natural' diagonal.

OPERATIONS PRINCIPLE

Process types indicate the position of processes on the volume–variety spectrum.

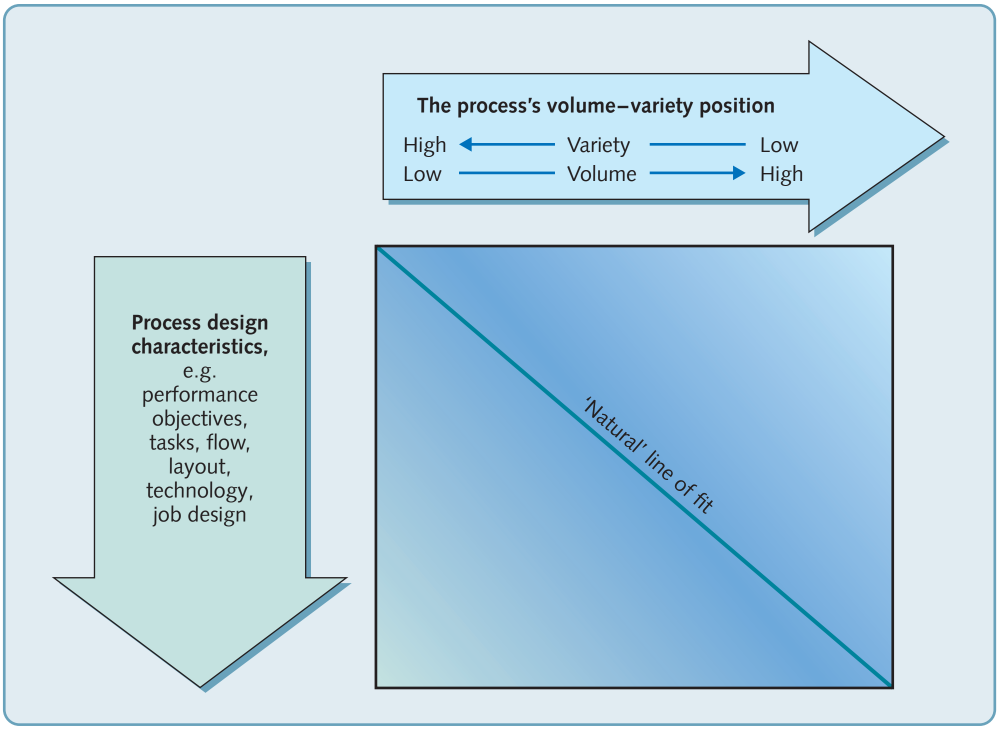

Figure 5.3 The elements of process design are strongly influenced by the volume–variety requirements placed on the process

Process types

Processes that inhabit different points on the diagonal of the product–process matrix are sometimes referred to as 'process types'. Each process type implies differences in the set of tasks performed by the process and in the way materials information, or customers flow through the process. Different terms are sometimes used to identify process types depending on whether they are predominantly manufacturing or service processes and there is some variation in how the names are used. This is especially so in service process types. It is not uncommon to find manufacturing terms used also to describe service processes. Perhaps most importantly, there is some degree of overlap between process types. The different process types are shown in Figure 5.4.

Project processes

Project processes are those that deal with discrete, usually highly customised products. Often the timescale of making the product is relatively long, as is the interval between the completions of each product. The activities involved in the process may be ill-defined and uncertain, sometimes changing during the process itself. Examples include advertising agencies, shipbuilding, most construction companies, and movie production companies, drilling oil wells and installing computer systems. Any process map for project processes will almost certainly be complex, partly because each unit of output is usually large with many activities occurring at the same time, and partly because the activities often involve significant discretion to act according to professional judgement. In fact, a process map for a whole project would be extremely complex, so rarely would a project be mapped, but small parts may be.

Jobbing processes

Jobbing processes also deal with very high variety and low volumes, but whereas in project processes each project has resources devoted more or less exclusively to it, in jobbing processes each 'product' has to share the operation's resources with many others. The process

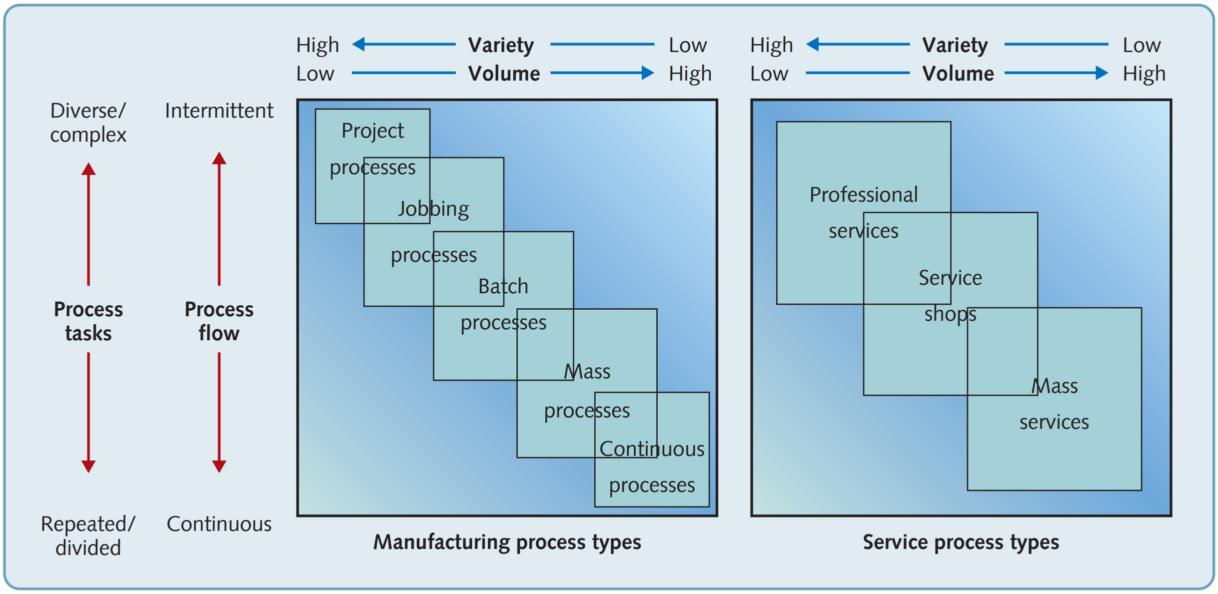

Figure 5.4 Different process types imply different volume–variety characteristics for the process

will work on a series of products but, although all the products will require the same kind of attention, each will differ in its exact needs. Examples of jobbing processes include many precision engineers such as specialist toolmakers, furniture restorers, 'make-to-measure' tailors and the printer who produces tickets for the local social event. Jobbing processes produce more and usually smaller items than project processes but, like project processes, the degree of repetition is low. Many jobs could be 'one-offs'. Again, any process map for a jobbing process could be relatively complex for similar reasons to project processes. Although jobbing processes sometimes involve considerable skill, they are usually more unpredictable than project processes.

Batch processes

Batch processes can look like jobbing processes, but without the degree of variety normally associated with jobbing. As the name implies, batch processes usually produce more than one 'product' at a time. So each part of the operation has periods when it is repeating itself, at least while the 'batch' is being processed. The size of the batch could be just two or three, in which case the batch process would differ little from jobbing, especially if each batch is a totally novel product. Conversely, if the batches are large, and especially if the products are familiar to the operation, batch processes can be fairly repetitive. Because of this, the batch type of process can be found over a wider range of volume and variety levels than other process types. Examples of batch processes include machine tool manufacturing, the production of some special gourmet frozen foods, the manufacture of most of the component parts that go into mass-produced assemblies such as automobiles and the production of most clothing. Batch process maps may look straightforward, especially if different products take similar routes through the process with relatively standard activities being performed at each stage.

Mass processes

Mass processes produce in high volume, usually with narrow effective variety. A car manufacturing plant, for example, might produce several thousand variants of cars if every option of engine size, colour and equipment is taken into account. Yet its effective variety is low because the different variants do not affect the basic process of production. The activities in the car manufacturing plant, like all mass processes, are essentially repetitive and largely predictable. In addition to the car plant, examples of mass processes include consumer durable manufacturers, most food processes such as a frozen pizza manufacturer, beer bottling plants and CD production. Process maps for this type of process will be straightforward sequences of activities.

Continuous processes

Continuous processes are one step beyond mass processes insomuch as they operate at even higher volume and often have even lower variety. Sometimes they are literally continuous in that their products are inseparable, being produced in an endless flow. Continuous processes are often associated with relatively inflexible, capital-intensive technologies with highly predictable flow. Examples of continuous processes include petrochemical refineries, electricity utilities, steel making and internet server farms. Like mass processes, process maps will show few elements of discretion, and although products may be stored during the process, the predominant characteristic of most continuous processes is of smooth flow from one part of the process to another.

Professional services

Professional services are high-variety, low-volume processes, where customers may spend a considerable time in the service process. Such services usually provide high levels of customisation, so contact staff are given considerable discretion. They tend to be people-based rather than equipment-based, with emphasis placed on the process (how the service is delivered) as much as the 'product' (what is delivered). Examples include management consultants, lawyers' practices, architects, doctors' surgeries, auditors, health and safety inspectors and some computer field service operations. Where process maps are used they are likely to be drawn predominantly at a high level. Consultants, for example, frequently use a predetermined set of broad stages, starting with understanding the real nature of the problem through to the implementation of their recommended solutions. This high-level process map guides the nature and sequence of the consultants' activities.

Service shops

Service shops are characterised by levels of customer contact, customisation, volume of customers and staff discretion that position them between the extremes of professional and mass services (see below). Service is provided via a mix of front- and back-office activities. Service shops include banks, High Street shops, holiday tour operators, car rental companies, schools, most restaurants, hotels and travel agents. For example, an equipment hire and sales organisation may have a range of equipment displayed in front-office outlets, while back-office operations look after purchasing and administration. The front-office staff have some technical training and can advise customers during the process of selling the product. Essentially, the customer is buying a fairly standardised 'product' but will be influenced by the process of the sale which is customised to the individual customer's needs.

Mass services

Mass services have many customer transactions and little customisation. Such services are often predominantly equipment based and 'product' oriented, with most value added in the back-office, sometimes with comparatively little judgement needed by front-office staff, who may have a closely defined job and follow set procedures. Mass services include supermarkets, a national rail network, an airport and many call centres. For example, airlines move a large number of passengers on their networks. Passengers pick a journey from the range offered. The airline can advise passengers on the quickest or cheapest way to get from A to B, but they cannot 'customise' the service by putting on special flights for them.

Moving off the natural diagonal

A process lying on the natural diagonal of the matrix shown in Figure 5.3 will normally have lower operating costs than one with the same volume–variety position that lies off the diagonal. This is because the diagonal represents the most appropriate process design for any volume–variety position. Processes that are on the right of the 'natural' diagonal would normally be associated with lower volumes and higher variety. This means that they are likely to be more flexible than seems to be warranted by their actual volume–variety position. That is, they are not taking advantage of their ability to standardise their activities. Because of this, their costs are likely to be higher than they would be with a process that was closer to the diagonal. Conversely, processes that are on the left of the diagonal have adopted a position that would normally be used for higher volume and lower variety processes. Processes will therefore be 'over-standardised' and probably too inflexible for their volume–variety position. This lack of flexibility can also lead to high costs because the process will not be able to change from one activity to another as readily as a more flexible process. One note of caution regarding this idea: although logically coherent, it is a conceptual model rather than something that can be 'scaled'. Although it is intuitively obvious that deviating from the diagonal increases costs, the precise amount by which costs will increase is very difficult to determine. Nevertheless, a first step in examining the design of an existing process is to check if it is on the natural diagonal of the product–process matrix. The volume–variety position of the process may have changed without any corresponding change in its design. Alternatively, design changes may have been introduced without considering their suitability for the processes volume–variety position.

OPERATIONS PRINCIPLE

Moving off the 'natural diagonal' of the product–process matrix will incur excess cost.

EXAMPLE

Water meters

The 'meter installation' unit of a water utility company installed and repaired water meters. The nature of each installation job could vary significantly because the metering requirements of each customer varied and because meters had to be fitted into very different water pipe systems. When a customer requested an installation or repair a supervisor would survey the customer's water system and transfer the results of the survey to the installation team of skilled plumbers. An appointment would then be made for a plumber to visit the customer's location and install or repair the meter on the agreed appointment date. The company decided to install for free a new 'standard' remote-reading meter that would replace the wide range of existing meters and could be read automatically using the customer's telephone line. This would save meter reading costs. It also meant a significant increase in work for the unit and more skilled plumbing staff were recruited. The new meter was designed to make installation easier by including universal quick-fit joints that reduced pipe cutting and jointing during installation. As a pilot, it was also decided to prioritise those customers with the oldest meters and conduct trials of how the new meter worked in practice. All other aspects of the installation process were left as they were.

The pilot was not a success. Customers with older meters were distributed throughout the company's area, so staff could not service several customers in one area and had to travel relatively long distances between customers. Also, because customers had not initiated the visit themselves, they were more likely to have forgotten the appointment, in which case plumbers had to return to their base and try to find other work to do. The costs of installation were proving to be far higher than forecast and the plumbers were frustrated at the waste of their time and the now relatively standardised installation job. The company decided to change its process. Rather than replace the oldest meters that were spread around its region, it targeted smaller

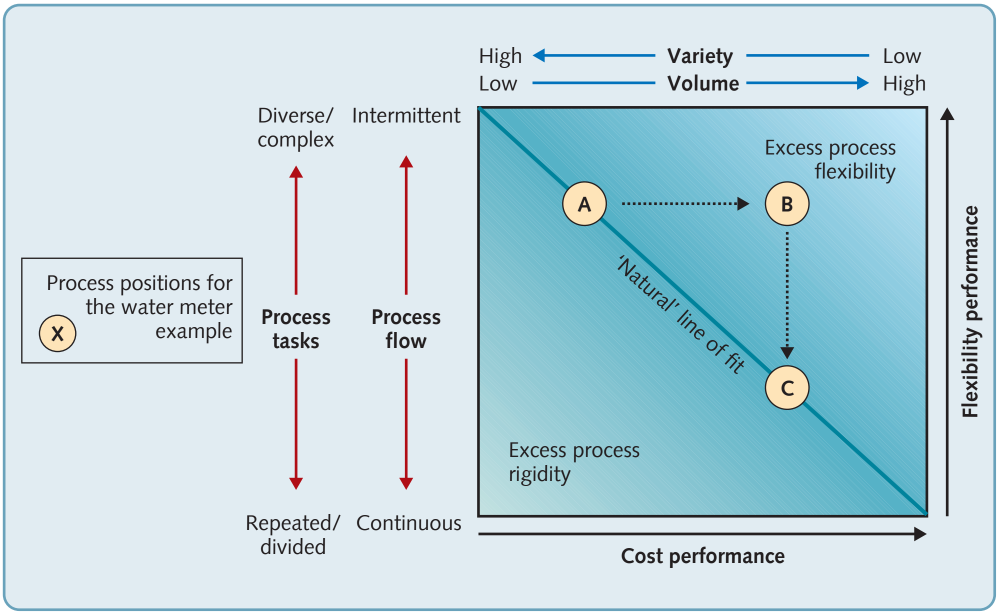

Figure 5.5 A product–process matrix with process positions from the water meter example

geographic areas to limit travelling time. It also cut out the survey stage of the process because, using the new meter, 98 per cent of installations could be fitted in one visit, minimising disruption to the customer and the number of missed appointments. Just as significantly, fully qualified plumbers were often not needed, so installation could be performed by less expensive labour.

This example is illustrated in Figure 5.5. The initial position of the installation process is at point A. The installation units were required to repair and install a wide variety of meters into a very wide variety of water systems. This needed a survey stage to assess the nature of the job and the use of skilled labour to cope with the complex tasks. The installation of the new type of meter changed the volume–variety position for the process by reducing the variety of the jobs tackled by the process and increasing the volume it had to cope with. However, the process was not changed. By choosing a wide geographic area to service, retaining the unnecessary survey stage and hiring over-skilled staff, the company was still defining itself as a high-variety, low-volume 'jobbing' process. The design of the process was appropriate for its old volume–variety position, but not the new one. In effect it had moved to point B in Figure 5.5. It was off the diagonal, with unnecessary flexibility and high operating costs. Redesigning the process to take advantage of the reduced variety and complexity of the job (position C on Figure 5.5) allowed installation to be performed far more efficiently.

Layout, technology and design

If movement down the natural diagonal of the product–process matrix changes the nature of a process, then the key elements of its design will also change. At this broad level, these 'key elements' of the design are the two 'ingredients' that make up process resources (technology, and people) and the way in which these ingredients are arranged within the process relative to each other. This latter aspect is usually called *layout*. In the remainder of the chapter, we start by discussing layout and then the design decisions that relate to process technology and the jobs that the people within the process undertake.

DIAGNOSTIC QUESTION

Are process layouts appropriate?

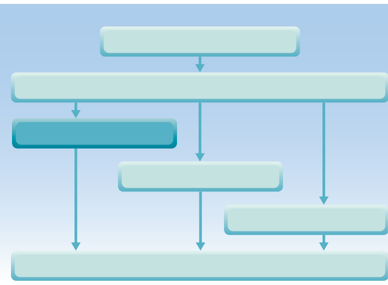

There is little point in having a well-sequenced process if in reality its activities are physically located in a way that involves excessive movement of materials, information or customers. Usually, the objective of the layout decision is to minimise movement, but, especially in an information-transforming process where distance is largely irrelevant, other criteria may dominate. For example, it may be more important to layout processes such that similar activities or resources are grouped together. So, an international bank may group its foreign exchange dealers together to encourage communication and discussion between them, even though the 'trades' they make are processed in an entirely different location. Some high-visibility processes may fix their layout to emphasise the behaviour of the customers who are being processed.

Layout should reflect volume and variety

Again, the layout of a process is determined partly by its volume and variety characteristics. When volume is very low and variety is relatively high, 'flow' may not be a major issue. For example, in telecommunications satellite manufacture each product is different, and because products 'flow' through the operation very infrequently, it is not worth arranging facilities to minimise the flow of parts through the operation. With higher volume and lower variety, flow becomes a far more important issue. If variety is still high, however, an entirely flow-dominated arrangement is difficult because there will be different flow patterns. For example, a library will arrange its different categories of books and its other services partly to minimise the average distance its customers have to 'flow' through the operation. But, because its customers' needs vary, it will arrange its layout to satisfy the majority of its customers (but perhaps inconvenience a minority). When the variety of products or services reduces to the point where a distinct 'category' with similar requirements becomes evident but variety is still not small, appropriate resources could be grouped into a separate cell. When variety is relatively small and volume is high, flow can become regularised and resources can be positioned to address the (similar) needs of the products or services, as in a classic flow line.

> OPERATIONS PRINCIPLE
> ***Resources in low-volume high-variety processes should be arranged to cope with irregular flow.***

> OPERATIONS PRINCIPLE
> ***Resources in high-volume low-variety processes should be arranged to cope with smooth, regular flow.***

Most practical layouts are derived from only four *basic layout types* that correspond to different positions on the volume–variety spectrum. These are illustrated diagrammatically in Figure 5.6 and are described below.

Fixed-position layout

Fixed-position layout is in some ways a contradiction in terms, since the transformed resources do not move between the transforming resources. Instead of materials, information or customers flowing through an operation, the recipient of the processing is stationary and the equipment, machinery, plant and people who do the processing move as necessary. This could be

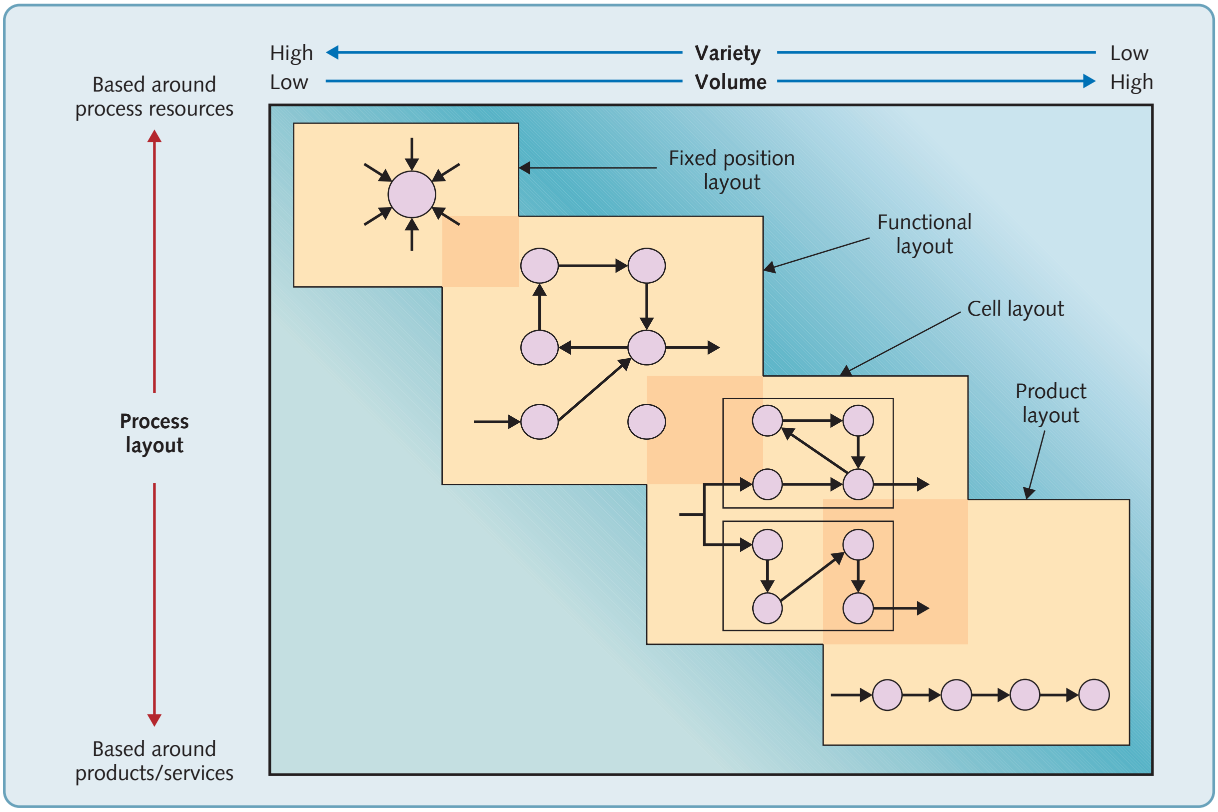

Figure 5.6 **Different process layouts are appropriate for different volume–variety combinations**

because the product or the recipient of the service is too large to be moved conveniently, or it might be too delicate to move, or perhaps it could object to being moved; for example:

- *Power generator construction* – the product is too large to move.
- *Open-heart surgery* – patients are too delicate to move.
- *High-class restaurant* – customers would object to being moved to where food is prepared.

Functional layout

Functional layout is so called because the functional needs and convenience of the transforming resources which constitute the processes dominate the layout decision. (Confusingly, functional layout can also be called 'process layout'). In functional layout, similar activities or resources (or those with similar needs) are located together. This may be because it is convenient to group them together, or that their utilisation can be improved. It means that when materials, information or customers flow through the operation, they will take a route from activity to activity according to their needs. Usually, this makes the flow pattern in the operation complex. Examples of process layouts include:

- *Hospitals* – some processes (e.g. radiography equipment and laboratories) are required by several types of patient.
- *Machining the parts for aircraft engines* – some processes (e.g. heat treatment) need specialist support (heat and fume extraction); some processes (e.g. machining centres) require the

same technical support from specialists; some processes (e.g. grinding machines) get high machine utilisation as all parts which need grinding pass through a single grinding section.

- *Supermarkets* – some products, such as tinned goods, are convenient to restock if grouped together. Some areas, such as those holding frozen vegetables, need the common technology of freezer cabinets. Others, such as the areas holding fresh vegetables, might be together because that way they can be made to look attractive to customers.

Cell layout

A cell layout is one where materials, information or customers entering the operation are preselected (or preselect themselves) to move to one part of the operation (or cell) in which all the transforming resources, to meet their immediate processing needs, are located. Internally, the cell itself may be arranged in any appropriate manner. After being processed in the cell, the transformed resources may go on to another cell. In effect, cell layout is an attempt to bring some order to the complexity of flow that characterises functional layout. Examples of cell layouts include:

- *Some computer component manufacture* – the processing and assembly of some types of computer parts may need a special area dedicated to producing parts for one particular customer who has special requirements such as particularly high-quality levels.
- *'Lunch' products area in a supermarket* – some customers use the supermarket just to purchase sandwiches, savoury snacks, cool drinks, etc. for their lunch. These products are often located close together in a 'cell' for the convenience of these customers.
- *Maternity unit in a hospital* – customers needing maternity attention are a well-defined group who can be treated together and who are unlikely to need the other facilities of the hospital at the same time that they need the maternity unit.

Product layout

Product layout involves locating people and equipment entirely for the convenience of the transformed resources. Each product, piece of information or customer follows a prearranged route in which the sequence of required activities corresponds to the sequence in which the processes have been located. The transformed resources 'flow' along a 'line'. This is why this type of layout is sometimes called flow or line layout. Flow is clear, predictable and therefore relatively easy to control. It is the high-volume and standardised requirements of the product or service which allow product layouts. Examples of product layout include:

- *Car assembly* – almost all variants of the same model require the same sequence of processes.
- *Self-service cafeteria* – generally the sequence of customer requirements (starter, main course, dessert, drink) is common to all customers, but layout also helps control customer flow.

Layout selection

Getting the process layout right is important, if only because of the cost, difficulty and disruption of making any layout change. It is not an activity many businesses would want to repeat very often. Also, an inappropriate layout could mean that extra cost is incurred *every time* an item is processed. But more than this, an effective layout gives clarity and transparency to the flow of items through a process. There is no better way of emphasising that everyone's activities are really part of an overall process than by making the flow between activities evident to everyone.

One of the main influences on which type of layout will be appropriate, is the nature of the process itself, as summarised in its 'process type'. There is often some confusion between process types and layout types. Layout types are not the same as process types. Process types were described earlier in the chapter and indicate a broad approach to the organisation and operation of a process. Layout is a narrower concept but is very clearly linked to process type. Just as process type is governed by volume and variety, so is layout. But for any given process type there are usually at least two alternative layouts. Table 5.1 summarises the alternative layouts for particular process types. Which of these is selected, or whether some hybrid layout is chosen, depends on the relative importance of the performance objectives of the process, especially cost and flexibility. Table 5.2 summarises.

Table 5.1 Alternative layout types for each process type

<table>
<tr><th>Manufacturing Process Type</th><th colspan="2">Potential Layout Types</th><th>Service Process Type</th></tr>
<tr><td>Project</td><td>Fixed-Position layout
Functional layout</td><td rowspan="2">Fixed-Position layout
Functional layout
Cell layout</td><td rowspan="2">Professional Service</td></tr>
<tr><td rowspan="2">Jobbing</td><td rowspan="2">Functional layout
Cell layout</td></tr>
<tr><td rowspan="3">Functional layout
Cell layout</td><td rowspan="3">Service Shop</td></tr>
<tr><td>Batch</td><td>Functional layout
Cell layout</td></tr>
<tr><td rowspan="2">Mass</td><td rowspan="2">Cell layout
Product layout</td></tr>
<tr><td rowspan="2">Cell layout
Product layout</td><td rowspan="2">Mass Service</td></tr>
<tr><td>Continuous</td><td>Product layout</td></tr>
</table>

Table 5.2 The advantages and disadvantages of the basic layout types

	Advantages	*Disadvantages*
Fixed-position	• Very high mix and product flexibility • Product or customer not moved or disturbed • High variety of tasks for staff	• High unit costs • Scheduling of space and activities can be difficult • Can mean much movement of plant and staff
Functional	• High mix and product flexibility • Relatively robust in the case of disruptions • Relatively easy supervision of equipment or plant	• Low facilities utilisation • Can have very high work-in-progress or customer queueing • Complex flow can be difficult to control
Cell	• Can give a good compromise between cost and flexibility for relatively high-variety operations • Fast throughput • Group work can result in good motivation	• Can be costly to rearrange existing layout • Can need more plant and equipment • Can give lower plant utilisation
Product	• Low unit costs for high volume • Gives opportunities for specialisation of equipment • Materials or customer movement are convenient	• Can have low mix flexibility • Not very robust if there is disruption • Work can be very repetitive

Layout and 'servicescapes'

In high-visibility processes (where the customer 'sees' much of the value-adding processing, see Chapter 1) the general look and feel of the process could be as important, if not more important, than more 'objective' criteria such as cost or distance moved. The term that is often used to describe the look and feel of the environment within an operation or process is its 'servicescape'. There are many academic studies that have shown that the servicescape of an operation plays an important role, both positive and negative, in shaping customers' views of how a high-visibility process is judged. The general idea is that ambient conditions, space factors and signs and symbols in a service operation will create an 'environment experience' both for employees and customers; and this environmental experience should support the service concept. The individual factors that influence this experience will then lead to certain responses (again, in both employees and customers). These responses can be put into three main categories:

- Cognitive (what people think)
- Emotional (what they feel) and
- Physiological (their body's experience).

However, remember that a servicescape will contain, not only objective, measureable and controllable stimuli, but also subjective, immeasurable and often uncontrollable stimuli, which will influence customer behaviour. The obvious example is other customers frequenting an operation. As well as controllable stimuli such as the colour, lighting, design, space and music, the number, demographics and appearance of one's fellow customers will also shape the impression of the operation.

OPERATIONS PRINCIPLE

Layout should include consideration of the look and feel of the operation to customers and/or staff.

EXAMPLE

Servicescapes in banks and factories[3]

Although retail banks are closing many of their branches, they are keen to make sure that those that do remain have a warm and friendly feeling. This is why they pay so much attention to the décor and design of their branches. This even extends to their smell. When customers visit one South London branch of Lloyds Bank they are greeted by an attractive aroma. It is the fragrance of 'white tea and thyme', a scent carefully chosen to conjure up the right feeling in the bank's customers. The building's designer, Sarah Harrison said, 'It gives that inviting feel, that welcoming feel. You can smell it on the High Street when the wind's blowing in the right direction.' But Lloyds is not the first bank to experiment with the opportunities of 'designing' sensory experience. Banks in some markets are offering customers freshly brewed coffee, partly because the aroma of coffee seems to lure potential new customers.

But do not assume that the idea of servicescapes applies only to service operations. Volkswagen's (VW) 'transparent factory' in Dresden, Germany, does not look like a traditional car assembly plant. Inside the factory that makes the up-market Phaeton sedan, the floors are expensive Canadian maple, the factory walls are made of clear glass (a loudspeaker outside imitates territorial bird sounds to keep birds from flying into the glass), and the workers all wear white coats and gloves. The whole operation has the atmosphere of a research laboratory rather than a factory. Partly, this is because the dirtier, noisier processes take place in another facility. Partly though, it is because the facility is as much a customer relations and marketing device as it is a production plant. Thousands of visitors visit the plant for tours each year. Its layout is visitor-friendly and is designed to receive 250 tourists per day by advance reservation. The ground floor houses a restaurant, and on the lower level there is a simulator that provides visitors with a virtual test drive of the Phaeton. Yet the Transparent Factory is also a serious manufacturing operation; it produces an average of 44 Phaetons a day, most of which are destined for China, Germany and South Korea.

DIAGNOSTIC QUESTION

Are process technologies appropriate?

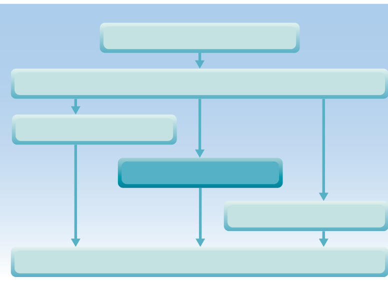

Process technologies are the machines, equipment and devices that help processes 'transform' materials, information and customers. It is a particularly important issue because few operations have been unaffected by the advances in process technology over the last two decades. And the pace of technological development is not slowing down. But, it is important to distinguish between *process technology* (the machines and devices that help to *create* products and services) and *product technology* (the technology that is embedded within the product or service and creates its specification or functionality). Some process technology, although not used for the actual creation of goods and services, nonetheless plays a key role in *facilitating* their creation. For example, the information technology systems that run planning and control activities can be used to help managers and operators run the processes. Sometimes this type of technology is called *indirect* process technology, and it is becoming increasingly important. Many businesses spend more on the computer systems that run their processes than they do on the direct process technology that creates their products and services.

Process technology should reflect volume and variety

Again, different process technologies will be appropriate for different parts of the volume–variety continuum. High-variety low-volume processes generally require process technology that is *general purpose,* because it can perform the wide range of processing activities that high variety demands. High-volume low-variety processes can use technology that is more *dedicated* to its narrower range of processing requirements. Within the spectrum, from general-purpose to dedicated process technologies, three dimensions in particular tend to vary with volume and variety. The first is the extent to which the process technology carries out activities or makes decisions for itself; that is, its degree of 'automation'. The second is the capacity of the technology to process work; that is, its 'scale' or 'scalability'. The third is the extent to which it is integrated with other technologies; that is, its degree of 'coupling' or 'connectivity'. Figure 5.7 illustrates these three dimensions of process technology.

OPERATIONS PRINCIPLE

Process technology in high-volume low-variety processes is relatively automated, large-scale and closely coupled when compared to that in low-volume, high-variety processes.

The degree of automation of the technology

To some extent, all technology needs human intervention. It may be minimal, for example the periodic maintenance interventions in a petrochemical refinery. Conversely, the person who operates the technology may be the entire 'brains' of the process, for example the surgeon using keyhole surgery techniques. Generally, processes that have high variety and low volume will employ process technology with lower degrees of automation than those with higher volume and lower variety. For example, investment banks trade in highly complex and sophisticated financial 'derivatives', often customised to the needs of individual clients, and each may be worth millions of dollars. The back-office of the bank has to process these deals to make

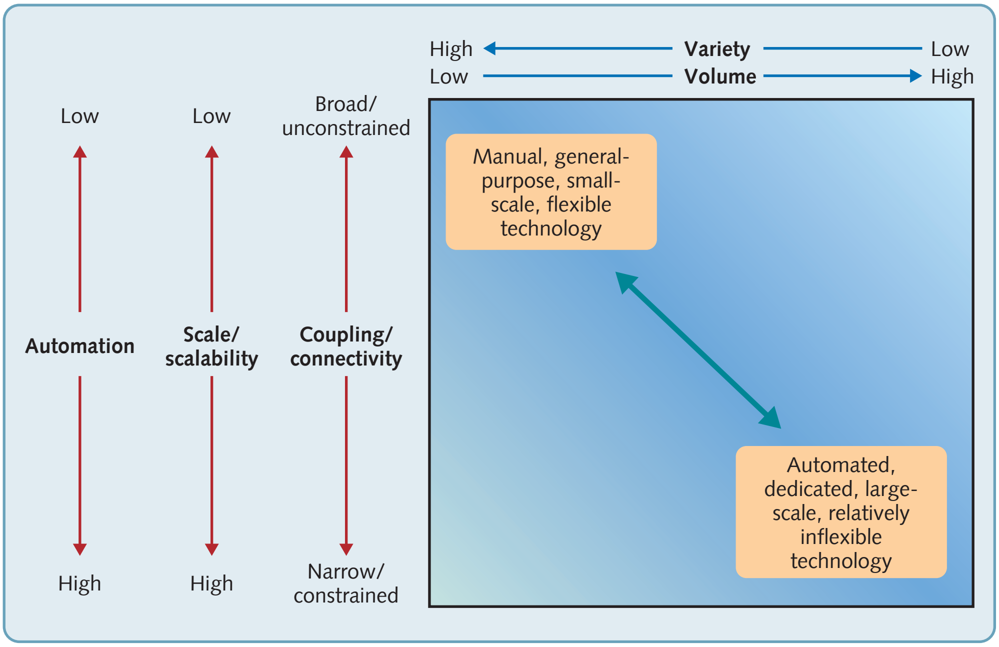

Figure 5.7 Different process technologies are important for different volume–variety combinations

sure that payments are made on time, documents are exchanged, and so on. Much of this processing will be done using relatively general-purpose technology such as spreadsheets. Skilled back-office staff make the decisions rather than the technology. Contrast this with higher volume, low-variety products, such as straightforward equity (stock) trades. Most of these products are simple and straightforward and are processed in very high volume of several thousand per day by 'automated' technology.

The scale/scalability of the technology

There is usually some discretion as to the scale of individual units of technology. For example, the duplicating department of a large office complex may decide to invest in a single, very large, fast copier, or alternatively in several smaller, slower copiers distributed around the operation's various processes. An airline may purchase one or two wide-bodied aircraft or a larger number of smaller aircraft. The advantage of large-scale technologies is that they can usually process items cheaper than small-scale technologies, but usually need high volume and can cope only with low variety. By contrast, the virtues of smaller-scale technology are often the nimbleness and flexibility that are suited to high-variety, lower volume processing. For example, four small machines can between them produce four different products simultaneously (albeit slowly), whereas a single large machine with four times the output can produce only one product at a time (albeit faster). Small-scale technologies are also more robust. Suppose the choice is between three small machines and one larger one. In the first case, if one machine breaks down, a third of the capacity is lost, but in the second, capacity is reduced to zero.

The equivalent to scale for some types of information processing technology is *scalability*. By scalability we mean the ability to shift to a different level of useful capacity quickly, and cost-effectively. Scalability is similar to absolute scale in so much as it is influenced by the same volume–variety characteristics. IT scalability relies on consistent IT platform architecture and the high process standardisation that is usually associated with high-volume and low-variety operations.

The coupling/connectivity of the technology

Coupling means the linking together of separate activities within a single piece of process technology to form an interconnected processing system. Tight coupling usually gives fast process throughput. For example, in an automated manufacturing system products flow quickly without delays between stages, and inventory will be lower – it can't accumulate when there are no 'gaps' between activities. Tight coupling also means that flow is simple and predictable, making it easier to keep track of parts when they pass through fewer stages, or information when it is automatically distributed to all parts of an information network. However, closely coupled technology can be both expensive (each connection may require capital costs) and vulnerable (a failure in one part of an interconnected system can affect the whole system). The fully integrated manufacturing system constrains parts to flow in a predetermined manner, making it difficult to accommodate products with very different processing requirements. So, coupling is generally more suited to relatively low variety and high volume. Higher variety processing generally requires a more open and unconstrained level of coupling because different products and services will require a wider range of processing activities

DIAGNOSTIC QUESTION

Are job designs appropriate?

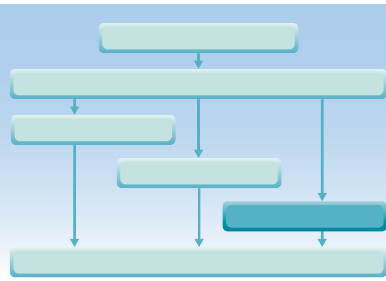

Job design is about how people carry out their tasks within a process. It defines the way they go about their working lives. It positions the expectations of what is required of them, and it influences their perceptions of how they contribute to the organisation. It also defines their activities in relation to their work colleagues and it channels the flows of communication between different parts of the operation. But, of most importance, it helps to develop the culture of the organisation: its shared values, beliefs and assumptions. Inappropriately designed jobs can destroy the potential of a process to fulfil its objectives, no matter how appropriate its layout or process technology. So jobs must be designed to fit the nature of the process. However, before considering this, it is important to accept that some aspects of job design are common to all processes, irrespective of what they do or how they do it. Consider the following.

- *Safety.* The primary and universal objective of job design is to ensure that all staff performing any task within a process are protected against the possibility of physical or mental harm.
- *Ethical issues.* No individual should be asked to perform any task that is either illegal or (within limits) conflicts with strongly held ethical beliefs.
- *Work/life balance.* All jobs should be structured so as to promote a healthy balance between time spent at work and time away from work.

Note that all these objectives of job design are also likely to improve overall process performance. However, the imperative to follow such objectives for their own sake transcends conventional criteria.

Job design should reflect volume and variety

As with other aspects of process design, the nature and challenges of job design are governed largely by the volume–variety characteristics of a process. An architect designing major construction projects will perform a wide range of very different, often creative and complex tasks, many of which are not defined at the start of the process, and most of which have the potential to give the architect significant job satisfaction. By contrast, someone in the architect's accounts office keying in invoice details has a job that is repetitive, has little variation, is tightly defined and is one that cannot rely on the intrinsic interest of the task itself to maintain job commitment. These two jobs will have different characteristics because they are part of processes with different volume and variety positions. Three aspects of job design in particular are affected by the volume–variety characteristics of a process: how tasks are to be allocated to each person in the process; the degree of job definition; and the methods used to maintain job commitment. Figure 5.8 illustrates this.

> OPERATIONS PRINCIPLE
>
> ***Job designs in high-volume low-variety processes are relatively closely defined with little decision-making discretion and needing action to help commitment when compared to those in low-volume, high-variety processes.***

The division of labour – how should tasks be allocated?

The most obvious aspect of any individual's job is how big it is; that is, how many of the tasks within any process are allocated to an individual. Should a single individual perform all the process? Alternatively, should separate individuals or teams perform each task? Separating tasks into smaller parts between individuals is called the *division of labour.* Perhaps its epitome is the assembly line, where products move along a single path and are built up by operators continually repeating a single task. This is the predominant model of job design in most high-volume

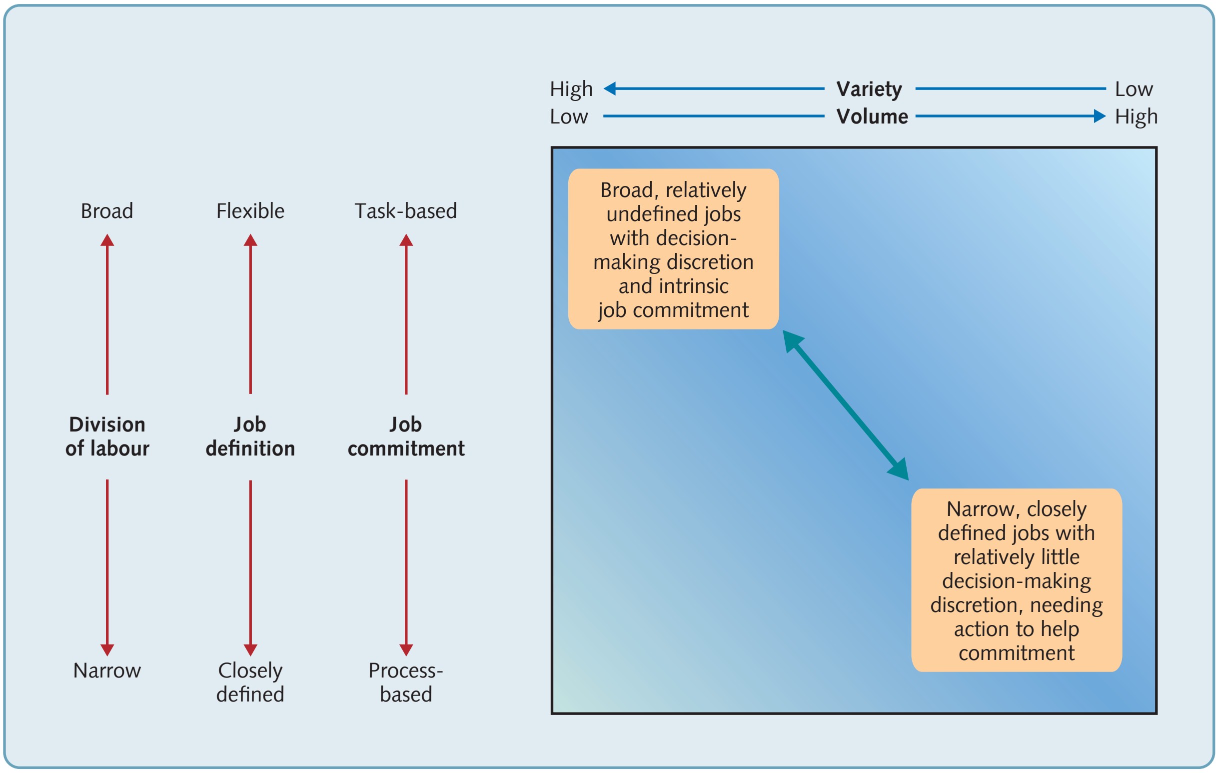

Figure 5.8 Different job designs are appropriate for different volume–variety combinations

low-variety processes. For such processes there are some *real advantages* in division-of-labour principles:

- *It promotes faster learning.* It is obviously easier to learn how to do a relatively short and simple task than a long and complex one, so new members of staff can be quickly trained and assigned to their tasks.
- *Automation becomes easier.* Substituting technology for labour is considerably easier for short and simple tasks than for long and complex ones.
- *Non-productive work is reduced.* In large, complex tasks the proportion of time between individual value-adding elements can be very high; for example in manufacturing, picking up tools and materials, putting them down again and generally searching and positioning.

There are also *serious drawbacks* to highly divided jobs:

- *It is monotonous.* Repeating the same task 8 hours a day and 5 days a week, is not fulfilling. This may lead to an increased likelihood of absenteeism, staff turnover and error rates.
- *It can cause physical injury.* The continued repetition of a very narrow range of movements can, in extreme cases, lead to physical injury. The over-use of some parts of the body (especially the arms, hands and wrists) can result in pain and a reduction in physical capability, called repetitive strain injury (RSI).
- *It can mean low flexibility*. Dividing a task up into many small parts often gives the job design a rigidity which is difficult to change under changing circumstances. For example, if an assembly line has been designed to make one particular product but then has to change to manufacture a quite different product, the whole line will need redesigning. This will probably involve changing every operator's set of tasks.
- *It can mean poor robustness.* Highly divided jobs imply items passing between several stages. If one of these stages is not working correctly, for example because some equipment is faulty, the whole operation is affected. On the other hand, if each person is performing the whole of the job, any problems will only affect that one person's output.

To what degree should jobs be defined?

Jobs in high-variety processes are difficult to define in anything but the most general terms. Such jobs may require tacit knowledge gained over time and through experience and often require individuals to exercise significant discretion in what they do and how they do it. Some degree of job definition is usually possible and advisable, but it may be stated in terms of the 'outcome' from the task rather than in terms of the activities within the task. For example, the architect's job may be defined in terms of *'achieving overall coordination, taking responsibility for articulating the overall vision of the project, ensuring stakeholders are comfortable with the process, etc.'* By contrast, a process with less variety and higher volume is likely to be defined more closely, with the exact nature of each activity defined and individual staff trained to follow a job step by step.

How should job commitment be encouraged?

Many factors may influence job commitment. An individual's job history and expectations, relationships with co-workers, personal circumstances, can all be important. So too are the volume and variety characteristics of the process, by defining the possible ways in which commitment can be enhanced. In high-variety processes, especially those with a high degree of staff discretion, job commitment is likely to come from the *intrinsic nature of the task* itself. Exercising skill and decision-making, for example, can bring their own satisfaction. Of course, commitment can be enhanced through extra responsibility, flexibility in working times, and so on, but the main motivator is the job itself. By contrast, low- variety high-volume jobs, especially those designed with a high division of labour and little discretion, can be highly alienating. Such jobs have relatively little intrinsic task satisfaction. It has to be *'designed into' the process* by emphasising the satisfaction to be gained from the performance of the process overall. A number of job design approaches have been suggested for achieving this in processes involving relatively repetitive work.

Job enlargement

Job enlargement involves allocating a larger number of tasks to individuals, usually by combining tasks that are broadly of the same type as those in the original job. This may not involve more demanding or fulfilling tasks, but it may provide a more complete and therefore slightly more meaningful job. If nothing else, people performing an enlarged job will not repeat themselves as often. For example, suppose that the manufacture of a product has traditionally been split up on an assembly line basis into ten equal and sequential jobs. If that job is then redesigned so as to form two parallel assembly lines of five people, each operator would have twice the number of tasks to perform.

Job enrichment

Like job enlargement, job enrichment increases the number of tasks in a job, but also implies allocating tasks that involve more decision-making, or greater autonomy, and therefore greater control over the job. This could include the maintenance of, and adjustments to, any process technology used, the planning and control of activities within the job, or the monitoring of quality levels. The effect is both to reduce repetition in the job *and* to increase personal development opportunities. So, in the assembly line example, each operator working on a job could also be allocated responsibility for carrying out routine maintenance and such tasks as record-keeping and managing the supply of materials.

Job rotation

Job rotation means moving individuals periodically between different sets of tasks to provide some variety in their activities. When successful, job rotation can increase skill flexibility and make a small contribution to reducing monotony. However, it is not always viewed as beneficial either by management (because it can disrupt the smooth flow of work) or by the people performing the jobs (because it can interfere with their rhythm of work).

Empowerment

Empowerment means enhancing individuals' ability, and sometimes authority, to change how they do their jobs. Some technologically constrained processes, such as those in chemical plants may limit the extent that staff can dilute their highly standardised task methods without consultation. Other less defined processes to empowerment may go much further.

Team-working

Closely linked to empowerment, team-based work organisation (sometimes called self-managed work teams), is where staff, often with overlapping skills, collectively perform a defined task and have some discretion over how they perform the task. The team may control such things as task allocation between members, scheduling work, quality measurement and improvement, and sometimes even the hiring of staff. Groups are described as 'teams' when the virtues of working together are being emphasised and a shared set of objectives and responsibilities is assumed.

Critical commentary

There could be three sets of criticisms prompted by the material covered in this chapter. The first relates to the separation of process design into two parts – positioning and analysis. It can reasonably be argued that this separation is artificial in so much as (as is admitted at the beginning of this chapter) the two approaches are very much interrelated. An alternative way of thinking about the topic would be to consider all aspects of the arrangement of resources together. This would include the issues of layout that have been discussed in this chapter, together with the more detailed process mapping issues describes in Chapter 6. The second criticism would challenge the core assumption of the chapter – that many significant process design decisions are influenced primarily by volume and variety. Whereas it is conventional to relate layout and (to a slightly lesser extent) process technology to volume–variety positioning, it is less conventional to do so for issues of job design. Some would argue that the vast majority of job design decisions will not vary significantly with volume and variety. The final criticism is also related to job design. Some academics would argue that our treatment of job design is over-influenced by the discredited (in their eyes) principles of the 'scientific' management movement that grew into 'work study' and 'time and motion' management.

SUMMARY CHECKLIST

- ☐ Do processes match volume–variety requirements?
- ☐ Are 'process types' understood and do they match volume–variety requirements?
- ☐ Can processes be positioned on the 'diagonal' of the product–process matrix?
- ☐ Are the consequences of moving away from the 'diagonal' of the product–process matrix understood?
- ☐ Are the implications of choosing an appropriate layout, especially the balance between process flexibility and low processing costs, understood?
- ☐ Are the process layouts appropriate?
- ☐ Which of the four basic layout types that correspond to different positions on the volume–variety spectrum is appropriate for each process?
- ☐ Have 'soft' servicescape factors been considered?
- ☐ Is process technology appropriate?
- ☐ Is the effect of the three dimensions of process technology (the degree of automation, the scale/scalability, and the coupling/connectivity of the technology) understood?
- ☐ Are job designs appropriate?
- ☐ Does job design ensure the imperative to design jobs that are safe, ethical, and promote an adequate work/life balance?
- ☐ Is the extent of division of labour in each process appropriate for its volume–variety characteristics?
- ☐ Is the extent of job definition in each process appropriate for its volume–variety characteristics?
- ☐ Are job commitment mechanisms in each process appropriate for its volume–variety characteristics?

CASE STUDY

McPherson Charles Solicitors

Grace Whelan, Managing Partner of McPherson Charles, welcomed the three solicitors into the meeting room. She outlined the agenda, essentially their thoughts and input into the rolling three-year plan. McPherson Charles, based in Bristol in the West of England, had grown rapidly to be one of the biggest law firms in the region, with 21 partners and around 400 staff; it was an ambitious partnership aiming to maintain its impressive growth record. The firm was managed through 15 teams, each headed by a partner. The meeting was intended to be the first stage in 'Plans for the Future', a programme to improve the effectiveness of the firm's operations. The three partners attending the meeting with Grace were Simon Reece (Family law), Kate Hutchinson (Property) and Hazel Lewis (Litigation). Grace asked for ideas on what the firm should prioritise in order to improve its performance further.

Simon Reece kicked things off, 'I think the first thing we need to agree on is that, for a professional service firm like ourselves, the quality of our people will always be the most important issue. We need to be absolutely confident that our staff not only have the best possible understanding of their own branch of the law, but also have the necessary client-relationship skills to consolidate our business position with increasingly demanding clients.' Hazel was not so sure. 'Of course I agree that the quality of our staff is an important issue, but that has always been true. What is new is the help we can get from some serious investment in technology and software. Just getting our systems and processes right would, I am sure, save us a lot of time and effort, and of course, reduce our cost base.'

'I really don't think spending more money on technology is the answer Hazel.' Simon continued, 'We need more time to really understand our clients and being process and IT focused just doesn't work for us, we need another way of managing. The key is increasing revenue, not penny pinching about costs, and to do that we need to really concentrate on relationship skills. Family law is like walking through a minefield; you can easily offend clients who are, almost by definition, in a highly emotional state. I think we need to make sure that senior members of staff with experience of managing client relationships pass on their knowledge to those who are less experienced.'

'I disagree, Simon.' It was Kate now, 'Our clients really are increasingly cost conscious and if we don't deliver value for money, word will spread very fast and our business will dry up. Much of the time we over-engineer our services. Why should we use highly qualified and expensive lawyers for every single task? I am convinced that, with slick systems and enhanced training, non-qualified people could do much of our work.'

Grace knew that solicitors liked nothing better than disagreeing – it was what they did best – and she knew that this was going to be another long meeting.

In very simple terms, these are the type of activities that each team was engaged in.

Simon Reece, Family law

'We are called the "family law" team but basically what we do is to help people through the trauma of divorce, separation and break up. Our biggest "high value" clients come to us because of word-of-mouth recommendation. Last year we had 89 of these "high value" clients and they all valued the personal touch that we were able to give them, getting to know them well and spending time with them to understand the, often "hidden" aspects of their case. These interviews cannot be rushed. These clients tend to be wealthy people and we will often have to drop everything and go off half way round the world to meet and discuss their situation. There are no standard procedures, every client is different, and everyone has to be treated as an individual. So we have a team of individuals who rise to the challenge each time and give great service. Of course, not all our clients are the super-rich. About a third of our annual family law income comes from about 750 relatively routine divorce and counselling cases. This work is a lot less interesting and I try to make sure that all my team have a mix of interesting and routine work over the year. I encourage them to exercise and develop their professional judgement. They are empowered to deal with any issues themselves or call in myself or one of the more senior members of the team for advice if appropriate. It is important to give this kind of responsibility to them so that they see themselves as part of a team. We are also the only part of the firm that has adopted an open-plan office arrangement centred around our specialist library of family case law.'

Hazel Lewis, Litigation

'The litigation team provides a key service for our commercial client base. Our primary work consists of handling bulk collections of debt. The group has 17 clients, of which 5 provide 85 per cent of the total volume. We work closely with the accounts departments of the client companies and have developed a semi-automatic approach to debt collection. Staff input data received from their clients into the system; from that point everything progresses through a pre-defined process, letters are produced, queries responded to and eventually debts collected, ultimately through court proceedings if necessary. Work tends to come in batches from clients and varies according to economic conditions, time of year and client sales activities. At the moment things are fairly steady; we had 872 new cases last week. The details of each case are sent over by the client; our people input the data onto our screens and set up a standard diary system for sending letters out. Some people respond quickly to the first letter and often the case is closed within a week or so; other people ignore letters and eventually we initiate court proceedings. We know exactly what is required for court dealings and have a pretty good process to make sure all the right documentation is available on the day.'

Kate Hutchinson, Property

'We are really growing fast and are building up an excellent reputation locally for being fast, friendly and giving value for money. Most of our work is "domestic", acting for individuals buying or selling their home, or their second home. Each client is allocated to a solicitor who becomes their main point of contact. But, given that we can have up to a hundred domestic clients a week, most of the work is carried out by the rest of the team behind the scenes. There is a relatively standard process to domestic property sales and purchases and we think that now we are pretty efficient at managing these standard jobs. Our process has four stages: one dealing with land registry searches, one liaising with banks who are providing the mortgage finance, one who make sure surveys are completed and one section that finalises the whole process to completion. We believe that this degree of specialisation can help us achieve the efficiencies that are becoming important as the market gets more competitive. Increasingly, we are also getting more complex "special" jobs. These are things like "volume re-mortgage" arrangements; rather complex "one-off" jobs, where a mortgage lender transfers a complex set of loan assets to another lender. "Special" jobs are always more complex than the domestic work and sometimes there are times when fast completion is particularly important and that can throw us a bit. The firm has recently formed partnerships with two large speculative builders, so we are getting into special "plot sales". All these "specials" do involve a lot of work and can occupy several members of the team for a time. We are now getting up to 25 of these "specials" each week, and they can be somewhat different to our normal work but we try to follow roughly the same process with them as the normal domestic jobs.'

Are each team's processes appropriate?

Grace was concerned. The three teams obviously had to cope with very different volumes of work and variety of activities. It was also clear that each team had developed different approaches to managing their processes. The question that she needed to address was, whether each team's approach was appropriate for the demands placed upon it.

QUESTIONS

1 What are the individual 'services' offered by each of the three teams?

2 Where would you place each service in a scale that goes from relatively low volume, relatively high variety, to relatively high volume, relatively low variety?

3 How would you describe each team's process in terms of its layout, the technology (if any) it uses, and the jobs of its staff?

4 Use the above information to draw a 'product'–process matrix. What does it indicate?

APPLYING THE PRINCIPLES

Some of these exercises can be answered by reading the chapter. Others will require some general knowledge of business activity and some might require an element of investigation. **Model answers for the first two questions are available on the student companion website.** For model answers to the other questions in this section, please ask your tutor/lecturer.

1 Visit a branch of a retail bank and consider the following questions:

(a) What categories of service does the bank seem to offer?

(b) To what extent does the bank design separate processes for each of its types of service?

(c) What are the different process design objectives for each category of service?

2 A company that produces a wide range of specialist educational kits for 5–10 year olds is based in an industrial unit arranged in a simple layout of with six departments, each performing a separate task. The layout is shown in Figure 5.9, together with the results of an investigation of the flow of parts and products between each department. However, the company plans to revamp its product range.

'This new range will totally replace our existing products, and although I believe our existing layout is fine for the current product range, I think that we will need to reconfigure our layout when we make the transition to the new product range.' (COO of the company)

The estimate for the flow between the departments when the new product range is introduced is shown in Figure 5.10.

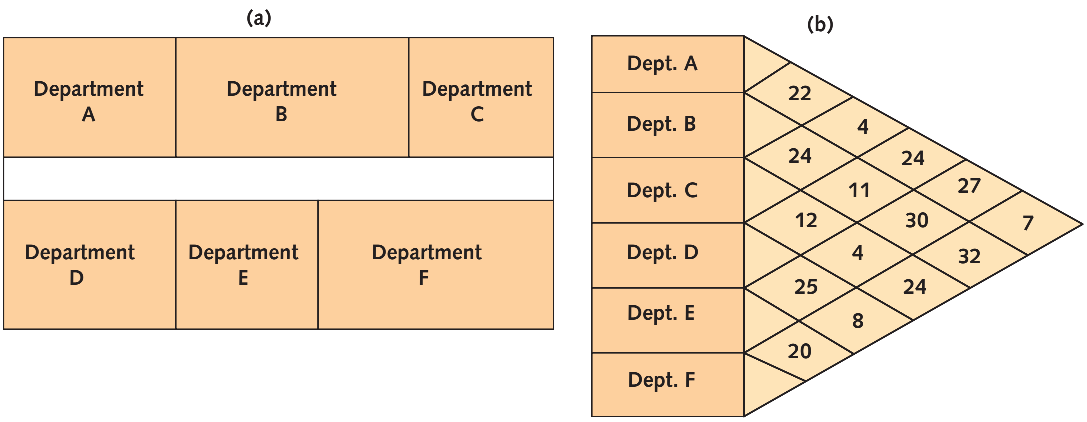

Figure 5.9 (a) the current layout of the educational kits producer, and (b) the current interdepartmental flow of parts and products (in pallet loads)

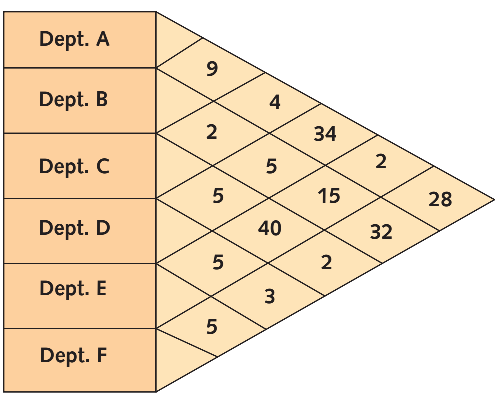

Figure 5.10 Estimated interdepartmental flow when the new product range is introduced

(a) Is the COO right in thinking that the current layout is right for the current product range?

(b) Assuming that the estimate of future interdepartmental flow is correct, how would you rearrange the factory?

3 The international Frozen Pizza Company (IFPC) operates in three markets globally. Market 1 is its largest market, where it sells 25,000 tons of pizza per year. In this market it trades under the name 'Aunt Bridget's Pizza' and positions itself as making pizza 'just as your Aunt Bridget used to make' (apparently she was good at it). It is also known for innovation, introducing new and seasonal pizza toppings on a regular basis. Typically, it sells around 20 varieties of pizza at any one time. Market 2 is smaller, selling around 20,000 tons per year under its 'Poppet's Pizza' brand. Although less innovative than Market 1, it still sells around 12 varieties of pizza. Market 3 is the smallest of the three, selling 10,000 tons per year of relatively high-quality pizzas under its 'Deluxe Pizza' brand. Like Aunt Bridget's Pizza, Deluxe Pizza also sells a relatively wide product range for the size of its market. Currently, both Markets 1 and 3 use relatively little automation and rely on high numbers of people, employed on a shift-system, to assemble their products. Market 2 has always been keen to adopt more automated production processes and uses a mixture of automated assembly and manual assembly. Now the management in Market 2 has developed an almost fully automated pizza assembly system (APAS). They claim that the APAS could reduce costs significantly and should be adopted by the other markets. Both Markets 1 and 3 are sceptical. 'It may be cheaper, but it can't cope with a high variety of products', is their response.

(a) Use the product–process matrix to explain the proposal by the management of Market 2.

4 A direct marketing company sells kitchen equipment through a network of local representatives working from home. Typically, individual orders usually contain 20 to 50 individual items. Much of the packing process is standardised and automatic. The Vice President of Distribution is proud of his distribution centre. 'We have a slick order fulfilment operation with lower costs per order, few packing errors, and fast throughput times. Our main problem is that the operation was designed for high volumes but the direct marketing business using representatives is, in general, on a slow but steady decline.' Increasingly, customers are moving towards using the company's recently launched website or just buying from supermarkets and discount stores. Bowing to the inevitable, the company has started selling its products through discount stores. The problem is, 'how to distribute their products through these new channels?' Should they modify their existing fulfilment operation or subcontract the business to specialist carriers? 'Although our system is great at what it does, it would be difficult to cope with very different types of order. Website orders will mean dealing with a far greater number of individual customers, each of whom will place relatively small orders for one or two items. We are not designed to cope with that kind of order. We would have the opposite problem delivering to discount stores. There, comparatively few customers would place large orders for a relatively narrow range of products. That is the type of job for a conventional distribution company. Another option would be to accept an offer from a non-competitor company who sell their products in a very similar way.'

(a) What are the implications of the different sales channels for the existing distribution centre?

5 Revisit the example at the beginning of this of the chapter that examines some of the principles behind supermarket layout. Then visit a supermarket and observe people's behaviour. You may wish to try and observe which areas they move slowly past and which areas they seem to move past without paying attention to the products. (You may have to exercise some discretion when doing this; people generally don't like to be stalked round the supermarket too obviously.)
Try and verify, as far as you can, some of the principles that were outlined in the box, and:

(a) What layout type is a conventional supermarket and how does it differ from a manufacturing operation using the same layout type?

(b) Some supermarkets are using customer tracking technology that traces the flow of customers through the shop. What are the benefits of using this type of technology for supermarkets?

Notes on chapter

1 Sources include: Sainsbury's press release (2015) 'Sainsbury's experiments with new supermarket design to make shopping quicker', 23 October; Usborne, S. (2012) 'The secrets of our supermarkets', the *Independent,* Friday 26 October.

2 Sources include: Knowles, T. (2017) 'Berkeley can see its prefabs sprouting up everywhere', *The Times,* 23 January; Davey, J. (2010) 'Today we built a house every hour', *The Sunday Times,* 31 January. Persimmon company website, Brown, G. (2010), 'Space4 growth gives boost to former car sector workers', *Birmingham Post,* 27 August.

3 Sources include: Milligan, B. (2016) 'Branch banking takes on a new fragrance for 2017', *BBC News,* 28 December. Also note that VW's transparent is due to close when the current model is phased out. It will be reconfigured to showcase the company's progress in the fields of 'electromobility and digitalization'.

TAKING IT FURTHER

Hammer, M. (1990) 'Reengineering Work: Don't Automate, Obliterate' , *Harvard Business Review, July–August.* *This is the paper that launched the whole idea of business processes and process management in general to a wider managerial audience. Slightly dated but worth reading.*

Harrington, H.J. (2011) *Streamlined Process Improvement*, *McGraw Hill Professional.*

Harvard Business Review (2011) *Improving Business Processes (Harvard Pocket Mentor), Harvard Business School Press.*

Hopp, W.J. and Spearman, M.L. (2011) *Factory Physics (*3rd edn*), Waveland Pr Inc.* *Very technical so don't bother with it if you aren't prepared to get into the maths. However, some fascinating analysis, especially concerning Little's Law.*

Jeston, J. and Nelis, J. (2014) *Business Process Management (3rd edn), Routledge.* *Written by two consultants, it gives chapter and verse on the orthodoxy of business process management; a broader topic than is covered in this chapter, but useful nonetheless.*

Panagacos, T. (2012) *The Ultimate Guide to Business Process Management: Everything you Need to Know and How to Apply it to Your Organization, CreateSpace Independent Publishing Platform.* *It is very much a practitioner book. It tells you how to do it.*

Ramaswamy, R. (1996) *Design and Management of Service Processes, Addison-Wesley Longman.* *A relatively technical approach to process design in a service environment.*

Process design 2 – Analysis

Introduction

The previous chapter set the broad parameters for process design; in particular it showed how volume and variety shape the resources of the process in terms of its appropriate layout, process technology and the design of its jobs. But this is only the beginning of process design. Within these broad resource parameters there are many, more detailed decisions regarding the activities carried out by the process that will dictate the way materials, information and customers flow through the process. Do not dismiss these detailed design decisions as merely the 'technicalities' of process design. They are important because they go a long way to determining the actual performance of the process in practice and eventually its contribution to the performance of the whole business. See Figure 6.1.

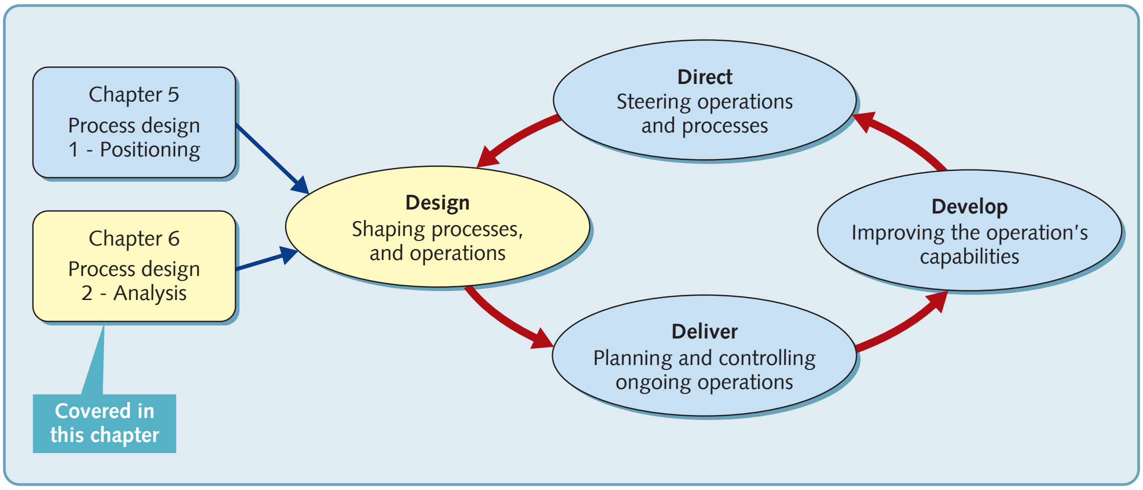

Figure 6.1 **Process design – analysis involves calculating the details of the process, in particular its objectives, sequence of activities, allocation of tasks and capacity, and its ability to incorporate the effects of variability**

EXECUTIVE SUMMARY

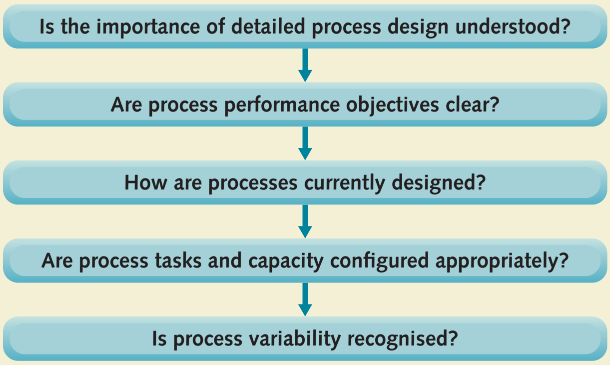

Is the importance of detailed process design understood?

How individual processes are designed is important to any operation because they are the building blocks of an operation. They link together to form 'internal process chains'. However, no chain is stronger than its weakest link, so the output from an operation will be limited by the effectiveness of its weakest process. This is why process design is so important. It involves calculating the details of the process, in particular its objectives, sequence of activities, allocation of tasks and capacity, and its ability to incorporate the effects of variability. It is the complementary activity to the broad positioning of processes that was described in the previous chapter.

Are process performance objectives clear?

The major objective of any process in the business is to support the business's overall objectives. Therefore process design must reflect the relative priority of the normal performance objectives, quality, speed, dependability, flexibility and cost. At a more detailed level, process design defines the way units flow through an operation. Therefore more 'micro' performance objectives are also useful in process design. Four in particular are used. These are throughput (or flow) rate, throughput time, the number of units in the process (work-in-process) and the utilisation of process resources.

How are processes currently designed?

Much process design is in fact redesign, and a useful starting point is to fully understand how the current process operates. The most effective way of doing this is to map the process in some way. This can be done at different levels using slightly different mapping techniques. Sometimes it is useful to define the degree of visibility for different parts of the process, Indicating how much of the process is transparent to customers.

Are process tasks and capacity configured appropriately?

This is a complex question with several distinct parts. First, it is necessary to understand the task precedence to be incorporated in the process. This defines what activities must occur before others. Second, it is necessary to examine how alternative process design options can incorporate series and parallel configuration. These are sometimes called 'long-thin' and 'short-fat' arrangements. Third, cycle time and process capacity must be calculated. This can help to allocate work evenly between the stages of the process (called balancing). Fourth, the relationship between throughput, cycle time and work-in-process must be established. This is done using a simple but extremely powerful relationship known as Little's Law (throughput time = work-in-process × cycle time).

Is process variability recognised?

In reality, processes have to cope with variability, both in terms of time and the tasks that are performed within the process. This variability can have very significant effects on process behaviour, usually to reduce process efficiency. Queuing theory can be used to understand this effect. In particular, it is important to understand the relationship between process utilisation and the number of units waiting to be processed (or throughput time).

DIAGNOSTIC QUESTION

Is the importance of detailed process design understood?

To 'design' is to conceive the looks, arrangement and workings of something *before it is constructed*. Process design should be treated at two levels – the broad, aggregated level and the more detailed level. The previous chapter took a broad approach by relating the resources in the process to the volume–variety position of the process. That will have identified the broad process type, and given some guidance as to the layout, process technology and job designs to be used within the process. This chapter takes a more detailed view of the activities undertaken by the process. However, in working out the details of a process design it is sometimes necessary to revisit the overall broad assumptions under which it is being designed. This is why the detailed analysis of process design covered in the chapter should always be thought through in the context of the broader process positioning issues covered in Chapter 5. The following two examples illustrate processes whose detailed design is important in determining their effectiveness.

How individual processes are designed is important to any operation because they are its building blocks. They link together to form 'internal process chains' through which flow material, information, or people. However, no chain is stronger than its weakest link. The output from a process is limited by the effectiveness of its worst stage and no operation is better than its weakest process. Configuring internal process networks involves designing each process and integrating individual processes to form an effective network. The root causes of business failure are often the result of the failures of everyday operational processes. And a prerequisite for long-term success is getting process design right.

EXAMPLE

Changi Airport[1]

Airports are complex operations. Really complex. Their processes handle passengers, aircraft, crew, baggage, commercial cargo, food, security, restaurants and numerous customer services. The operations managers, who oversee the daily operations of an airport, must cope with Aviation Administration rules and regulations, a huge number of airport service contracts, usually thousands of staff with a wide variety of specialisms, airlines with sometimes competing claims to service priority, customers some of whom fly every week and others who are a family of seven with two baby strollers that fly once a decade. Also their processes are vulnerable to disruptions from late arrivals, aircraft malfunction, weather, and the industrial action of workers two continents away, conflicts, terrorism and exploding volcanoes in Iceland. Designing the processes that can operate under these conditions must be one of the most challenging operations tasks. So to win prizes for 'Best Airport' customer

service and operating efficiency year after year has to be something of an achievement, which is what the sixth busiest international airport, Changi Airport in Singapore has done. As a major air hub in Asia, Changi serves more than 100 international airlines flying to some 300 cities in about 70 countries and territories worldwide. It handles almost 60 million passengers (that's roughly 10 times the size of Singapore's population). A flight takes off or lands at Changi roughly once every 90 seconds.

In 2017 Changi Airport opened its new Terminal 4 that was started in 2013. The new US$1.03 billion T4 is expected to handle about 16 million passengers per year and will increase the airport's annual passenger handling capacity to 82 million. Every stage of the customers' journey through the terminal has been designed to be as smooth as possible. Each stage in the customer journey must have enough capacity to cope with anticipated demand. A new overhead bridge will be built across the airport boulevard connecting T4 with Singapore's highway system and enable the movement of cars, buses and airside vehicles. Two new car parks will accommodate up to 1,500 vehicles. The terminal will be internally connected to the new car parks via sheltered links. Once passengers arrive at the two-storey terminal building they will pass through kiosks and automated options for self-check-in, self-bag-tagging and self-bag-drops. Their bags will then be transported to the aircraft via an advanced and automated baggage-handling system. Similarly, automated options will be offered at immigration counters and departure-gate boarding. Biometric technology and fast and seamless travel (FAST) services are being implemented at the terminal to speed passenger throughput, reduce manpower and increase efficiency. After security checks, passengers find themselves in 15,000 m^2 of shopping, dining, liquor, tobacco, perfumes, cosmetics and other retail spaces. It will implement a new walk-through retail concept. It will feature local, cultural and heritage-theme restaurants, as well as retail stores. This space also features a 300-m long Central Galleria, which will be a glazed open space that visually connects the departure, check-in, arrival and transit areas across the terminal. The emphasis on the aesthetic appeal of the terminal is something that Changi has long considered important. It already boasts a butterfly garden, orchid and sunflower gardens as well as a koi pond.

The feelings of passengers using the terminal are an important part of its design. Mr Yam Kum Weng, Executive Vice President of CAG, one of the companies helping to develop the design for the new terminal, said, '*T4 breaks new ground in passenger experience for travellers, while ensuring smooth and efficient operations for airlines and airport agencies. Architecturally, the design of T4 will be functional, and yet have its own distinct character compared to the other three terminals at Changi Airport. Our focus for the development of T4 will be on its interior and ensuring that the design and layout continues to be passenger-centric and user-friendly. It will offer what passengers want – a good range of leisure amenities, convenient facilities and attractive commercial offerings.*' And with so many different companies involved in the day-to-day operation of the airport it was vital to include as many stakeholders as possible during the design. Workshops were conducted with various stakeholders, including airlines, ground handlers, immigration and security agencies, retail and food and beverage operators as well as other users to ensure that the T4 design met the needs of each party. The objective was to ensure that T4, when operational, can deliver a seamless and refreshing experience for travellers, and also be a place where staff will feel proud and motivated to work.

EXAMPLE

Processes for even faster food[2]

The quick service restaurant (QSR) industry reckon that the very first drive-through dates back to 1928 when Royce Hailey first promoted the drive-through service at his Pig Stand restaurant in Los Angeles. Customers would simply drive by the back door of the restaurant where the chef would come out and deliver the restaurant's famous 'Barbequed Pig' sandwiches. Today, drive-through processes are slicker, faster. They are also more common; in 1975,

McDonald's did not have any drive-throughs, and now more than 90 per cent of its US restaurants incorporate a drive-through process and about 70 per cent of McDonald's orders come from people who never leave their car. Says one industry specialist, '*There are a growing number of customers for whom fast food is not fast enough. They want to cut waiting time to the very minimum without even getting out of their car. Meeting their needs depends on how smooth we can get the process.*'

The competition to design the fastest and most reliable drive-through process is fierce. Starbuck's drive-throughs have strategically placed cameras at the order boards so that servers can recognise regular customers and start making their order even before it's placed. Burger King has experimented with sophisticated sound systems, simpler menu boards and see-through food bags to ensure greater accuracy (no point in being fast if you don't deliver what the customer ordered). These details matter. McDonald's reckon that their sales increase one per cent for every six seconds saved at a drive-through. But there may be a trade-off in this type of process. McDonalds introduced a new process called 'ask, ask, tell', requiring employees to repeat the orders back to customers three times: when they place their order, at the window when they pay; and finally when they hand customers their food. It slows the process down marginally, but ensures that orders are right.

What do they have in common?

Both examples highlight a number of process design issues. The first is that the payback from good process design is clearly significant, or to reverse this logic, bad process design will quickly show in terms of cost and customer service. If you have been in an airport terminal when the baggage handling system has malfunctioned, or security checks are taking longer than normal, or when it's just not coping with a holiday rush, you know the frustration of being the victim of poorly designed processes. And although competition between airports is perhaps not as fierce as between fast food drive-throughs, in the long term poorly designed passenger, aircraft, cargo and baggage-handling process will adversely affect any airport's viability. Similarly, quick service restaurant operations devote time and effort to the design process, assessing the performance of alternative process designs in terms of efficiency, quality and, above all, throughput time. In this example, note also how it is difficult to separate the design of the process from the design of the product or service that it produces. The commonly offered 'combo meals' are designed with the constraints and capabilities of the drive-through process in mind. Both examples also illustrate the importance of not being afraid to analyse processes at a very detailed level. This may include thoroughly understanding current processes so that any improvement can be based on the reality of what happens in practice. It will certainly involve allocating the tasks and associated capacity very carefully to appropriate parts of the process. And, for most processes, it will also involve a design that is capable of taking into consideration the variability that exists in most human tasks. These are the topics covered in this chapter.

OPERATIONS PRINCIPLE

Processes should always be designed to reflect customer and/or market requirements.

> **OPERATIONS PRINCIPLE**
> *Not all processes are the same, someneed to be highly specified, but not all. Process simply means how we do things.*

The very idea of 'process design' can be controversial

The whole idea of detailed process design can polarise opinion among operations managers. Two very opposing views are often expressed. But both can be simplistic.

The first view is that everyone and everything is a process. And, if all work is a process, one should strive to apply a standardised set of disciplines and rules that can make what we do both more effective and more efficient. The objective, it is argued, is to find the best sequence of activities that describe the way things should be done, then preserve it in a formal set of rules that will ensure the same result 'every time, all the time'. The advantage of formal processes is the discipline it imposes. The opposing view believes that effective management must always rely on understanding, imagination and insight. You can't, according to this view, reduce everything to a standardised set of steps. People are the most important element in any processes, and people can't be treated like cogs in a machine. Processes, they say, describe everything as a machine-like set of mindless routine activities. Everything is about people, and effective management must always rely on the quality of people and how they are led with inspiration and imagination. Constraining all activities into a process straightjacket kills the essential humanity of working life.

Of course, both of these rather stereotyped views are extreme, but they do represent the two ends of a spectrum of attitudes commonly expressed. But think about what we mean by a process – a set of resources and activities that produce value. That is exactly what we mean here by a process – no more and no less. Yes, everyone is part of a process. But that does not imply that all processes should be highly specified so jobs are de-skilled, or that jobs should always be done in the same way. Some may need to be, but not all. The argument, sometimes referred to as 'people versus process', is based on a false distinction. People are the most important element in most processes, and 'process' is simply how we do things. In other words, not all processes are the same, nor should they be treated in the same way. The way processes are designed should depend on what we want them to do.

DIAGNOSTIC QUESTION

Are process performance objectives clear?

> **OPERATIONS PRINCIPLE**
> *Process performance can be judged in terms of the levels of quality, speed, dependability, flexibility and cost they achieve.*

The whole point of process design is to make sure that the performance of the process is appropriate for whatever it is trying to achieve. For example, if an operation competed primarily on its ability to respond quickly to customer requests, its processes would need to be designed to give fast throughput times. This would minimise the time between customers requesting a product or service and them receiving it. Similarly, if an operation competed on low price, cost-related objectives are likely to dominate its process design. Some kind of logic should link what the operation as a whole is attempting to achieve, and the performance objectives of its individual processes. This is illustrated in Table 6.1.

In addition to the standard performance objectives, process designs are increasingly being judged on their environmental impact. The following example of Ecover illustrates one organisation with an ethical philosophy that stresses environmental impact objectives.

Table 6.1 The impact of strategic performance objectives on process design objectives and performance

Strategic performance objective	*Typical process design objectives*	*Some benefits of good process design*
Quality	• Provide appropriate resources, capable of achieving the specification of product of services • Error-free processing	• Products and service produced 'on-specification' • Less recycling and wasted effort within the process
Speed	• Minimum throughput time • Output rate appropriate for demand	• Short customer waiting time • Low in-process inventory
Dependability	• Provide dependable process resources • Reliable process output timing and volume	• On-time deliveries of products and services • Less disruption, confusion and rescheduling within the process
Flexibility	• Provide resources with an appropriate range of capabilities • Change easily between processing states (what, how, or how much is being processed?)	• Ability to process a wide range of products and services • Low cost/fast product and service change • Low cost/fast volume and timing changes • Ability to cope with unexpected events (e.g. supply or a processing failure)
Cost	• Appropriate capacity to meet demand • Eliminate process waste in terms of: – excess capacity – excess process capability – in-process delays – in-process errors – inappropriate process inputs	• Low processing costs • Low resource costs (capital costs) • Low delay/inventory costs (working capital costs)

EXAMPLE

Ecover's ethical operation design[3]

Ecover is the world's largest producer of ecological detergents and eco-friendly cleaning and laundry products, with factories in France and Belgium that embody the company's commitment to sustainability. For example, Ecover factories operate entirely on green electricity – the type produced by wind generators, tidal generators and other natural sources. They also make the most of the energy they do use by choosing energy-efficient lighting. And, although the machinery they use in the factories is standard for the industry, they keep their energy and water consumption down by choosing low-speed appliances that don't require water to clean them. For example, the motors on their mixing machines can mix 25 tonnes of Ecover liquid while 'consuming no more electricity than a few flat irons'. And they have a 'squeezy gadget that's so efficient at getting every last drop of product out of the pipes, they don't need to be rinsed through'. Ecover say that they. . . .. 'hate waste, so we're big on recycling. We keep the amount of packaging used in our products to a minimum, and make sure that whatever cardboard or plastic we do use can be recycled, reused or refilled. It's an ongoing process of improvement; in fact, we've recently developed a new kind of green plastic we like to call 'Plant-astic' that's 100% renewable, reusable and recyclable – and made from sugarcane'. Even the building is ecological, designed to follow the movement of the sun so that production takes place with the maximum amount of natural daylight (saving power). The factory's frame is built from pine, using bricks that are made from clay, wood pulp and mineral waste, requiring less energy to bake. The factories' roofs are covered in a flowering plant that gives insulation all year round, so effective, that they don't need heating or air conditioning.

Process flow objectives

All the strategic performance objectives translate directly to process design as shown in Table 6.1. But, because processes will be managed at a very operational level, process design also needs to consider a more 'micro' and detailed set of objectives. These are largely concerned with flow through the process. When whatever is being 'processed' (we shall refer to these as 'units' irrespective of what they are) enter a process they will progress through a series of activities where they are 'transformed' in some way. Between these activities the units may dwell for some time in inventories, waiting to be transformed by the next activity. This means that the time that a unit spends in the process (its throughput time) will be longer than the sum of all the transforming activities that it passes through. Also the resources that perform the processes activities may not be used all the time because not all units will necessarily require the same activities and the capacity of each resource may not match the demand placed upon it. So neither the units moving through the process, nor the resources performing the activities may be fully utilised. Because of this, the way that units leave the process is unlikely to be exactly the same as the way they arrive at the process. Figure 6.2 illustrates some of the 'micro' performance flow objectives that describe process flow performance and the process design factors that influence them. The flow objectives are:

OPERATIONS PRINCIPLE

Process flow objectives should include throughput rate, throughput time, work-in-process and resource utilisation, all of which are interrelated.

- Throughput rate (or flow rate) is the rate at which units emerge from the process, i.e. the number of units passing through the process per unit of time.
- Throughput time is the average elapsed time taken for inputs to move through the process and become outputs.
- The number of units in the process (also called the 'work-in-process', or in-process inventory), as an average over a period of time.
- The utilisation of process resources is the proportion of available time that the resources within the process are performing useful work.

The design factors that will influence the flow objectives are the:

- variability of input arrival to the process
- configuration of the resources and activities within the process
- capacity of the resources at each point in the process and
- variability of the activities within the process.

As we examine each of these design factors, we will be using a number of terms that, although commonly used within process design, need some explanation. These terms will be described in the course of the chapter, but for reference, Table 6.2 summarises the terms.

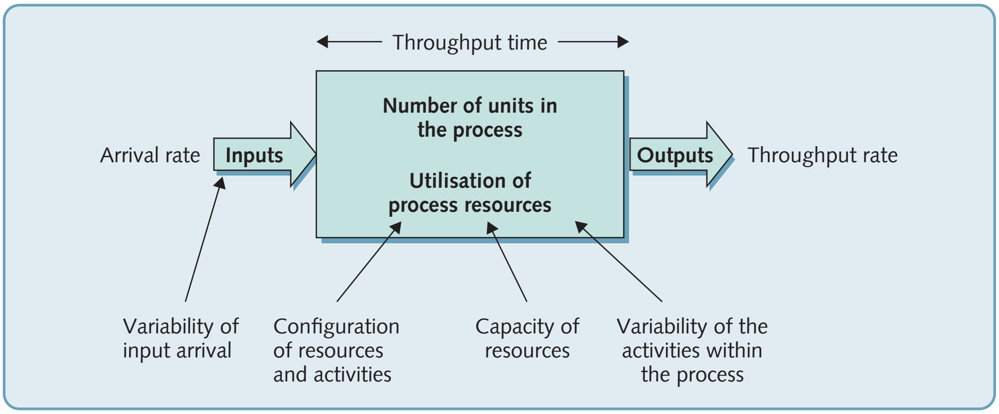

Figure 6.2 Micro process performance objectives and process design factors

Table 6.2 Some common process design terms

	Some common process design terms[4]
Process task	- The sum of all the activities that must be performed by the process.
Work content of the process task	- The total amount of work within the process task measured in time units.
Activity	- A discrete amount of work within the overall process task.
Work content of an activity	- The amount of work within an activity measured in time units.
Precedence relationship	- The relationship between activities expressed in terms of their dependencies, i.e. whether individual activities must be performance before other activities can be started.
Cycle time	- The average time that the process takes between completions of units.
Throughput rate	- The number of units completed by the process per unit of time (= 1/cycle time).
Process stage	- A work area within the process through which units flow; it may be responsible for performing several activities.
Bottleneck	- The capacity constraining stage in a process; it governs the output from the whole process.
Balancing	- The act of allocating activities as equally as possible between stages in the process.
Utilisation	- The proportion of available time that the process, or part of the process, spends performing useful work.
Starving	- Underutilisation of a stage within a process caused by inadequate supply from the previous stage.
Blocking	- The inability of a stage in the process to work because the inventory prior to the subsequent stage is full.
Throughput time	- The elapsed time between a unit entering the process and it leaving the process.
Queue time	- The time a unit spends waiting to be processed.

DIAGNOSTIC QUESTION

How are processes currently designed?

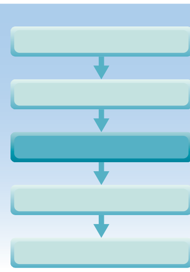

Existing processes are not always sufficiently well defined or described. Sometimes this is because they have developed over time without ever being formally recorded, or they may have been changed (perhaps improved) informally by the individuals who work in the process. But processes that are not formally defined can be interpreted in different ways, leading to confusion and inhibiting improvement. So, it is important to have some recorded visual descriptor of a process that can be agreed by all those who are involved in it. This is where process mapping comes in.

OPERATIONS PRINCIPLE
Process mapping is needed to expose the reality of process behaviour.

Process mapping

Process mapping (or process blueprinting as it is sometimes called) at its most basic level involves describing processes in terms of how the activities within the process relate to each other. There are many, broadly similar, techniques that can be used for process mapping. However, all the techniques have two main features:

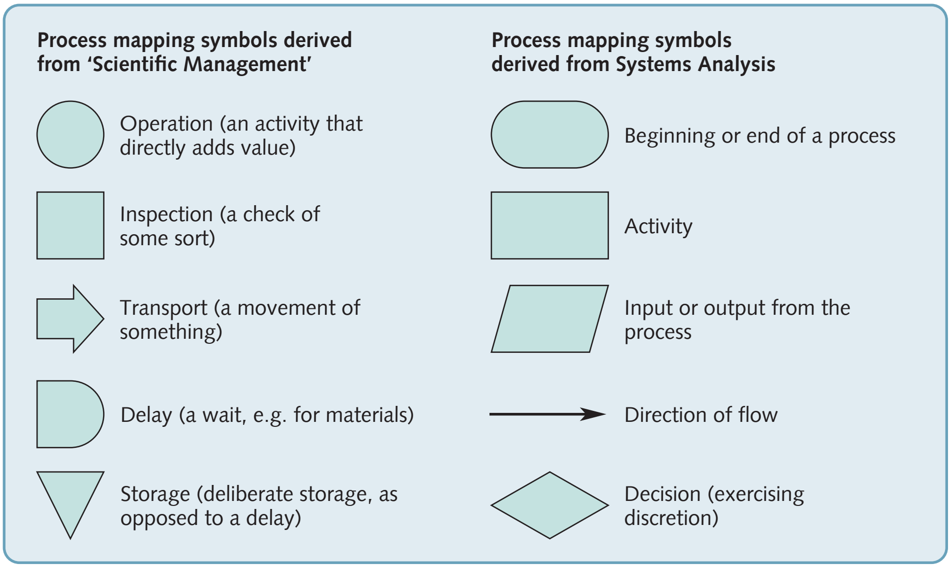

Figure 6.3 **Some common process mapping symbols**

1. They identify the different types of activity that take place during the process.
2. They show the flow of materials or people or information through the process (or, put another way, the sequence of activities that materials, people or information are subjected to).

Different process mapping symbols are sometimes used to represent different types of activity. They can be arranged in order, and in series or in parallel, to describe any process. And although there is no universal set of symbols used all over the world, some are relatively common. Most derive either from the early days of 'scientific' management around a century ago, or, more recently, from information system flowcharting. Figure 6.3 shows some of these symbols.

EXAMPLE

Theatre lighting operation

Figure 6.4 shows one of the processes used in a theatre lighting operation. The company hires out lighting and stage effects equipment to theatrical companies and event organisers. Customers' calls are routed to the store technician. After discussing their requirements, the technician checks the equipment availability file to see if the equipment can be supplied from the company's own stock on the required dates. If the equipment cannot be supplied in-house, customers may be asked whether they want the company to try and obtain it from other possible suppliers. This offer depends on how busy and how helpful individual technicians are. Sometimes customers decline the offer and a 'Guide to Customers' leaflet is sent to the customer. If the customer does want a search, the technician will call potential suppliers in an attempt to find available equipment. If this is not successful the customer is informed, but if suitable equipment is located it is reserved for delivery to the company's site. If equipment can be supplied from the company's own stores, it is reserved on the equipment availability file and the day before it is required a 'kit wagon' is taken to the store where all the required equipment is assembled, taken back to the workshop checked, and if any equipment is faulty it is repaired at this point. After that it is packed in special cases and delivered to the customer.

Different levels of process mapping

For a large process, drawing process maps at this level of detail can be complex. This is why processes are often mapped at a more aggregated level, called high-level process mapping, before more detailed maps are drawn. Figure 6.5 illustrates this for the total *'supply and install lighting'* process in the stage lighting operation. At the highest level the process can be drawn

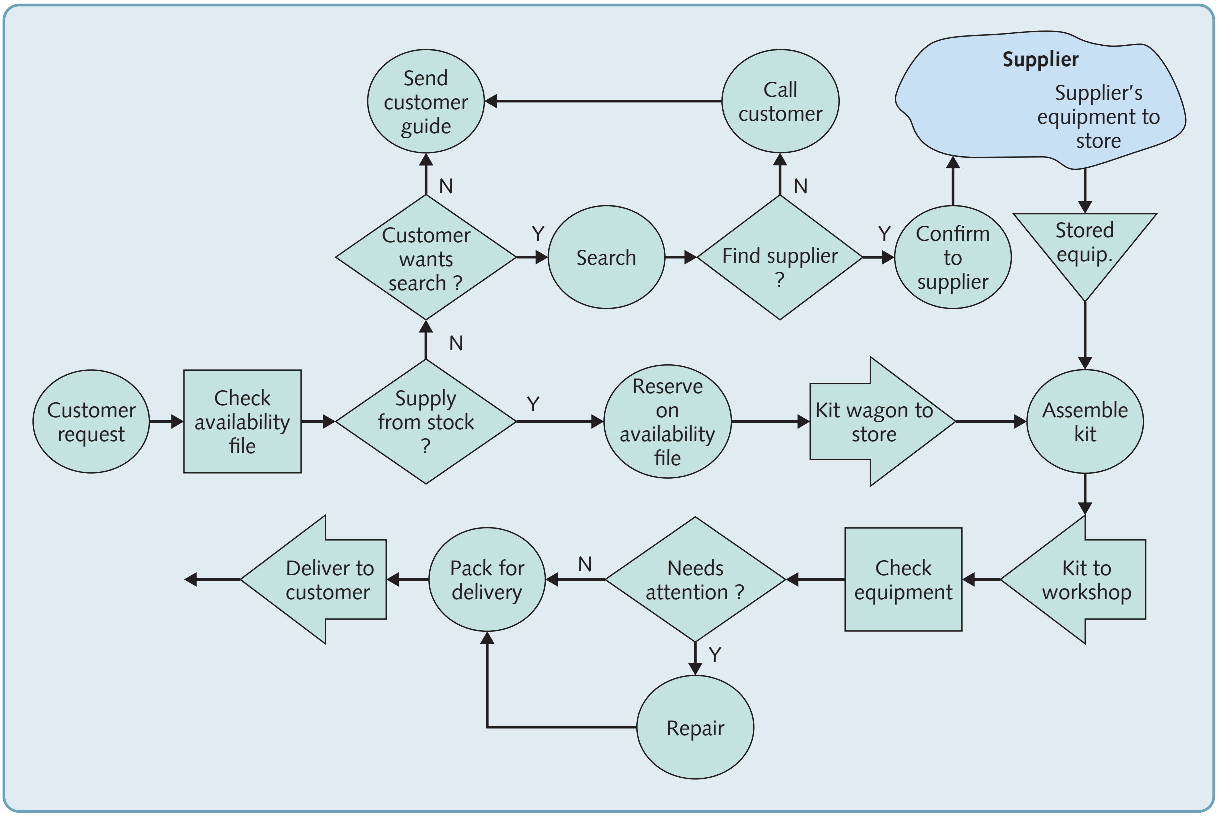

Figure 6.4 **Process map for 'enquire to delivery' process at stage lighting operation**

simply as an input–transformation–output process with materials and customers as its input resources and lighting services as outputs. No details of how inputs are transformed into outputs are included. At a slightly lower or more detailed level, what is sometimes called an outline process map (or chart) identifies the sequence of activities but only in a general way. So the process of '*enquire to delivery*' that is shown in detail in Figure 6.4 is here reduced to a single activity. At the more detailed level, all the activities are shown in a 'detailed process map' (the activities within the process 'install and test' are shown).

Although not shown in Figure 6.5 an even more micro set of process activities could be mapped within each of the detailed process activities. Such a micro detailed process map could specify every single motion involved in each activity. Some quick service restaurants, for example, do exactly that. In the lighting hire company example most activities would not be mapped in any more detail than that shown in Figure 6.5. Some activities, such as 'return to base', are probably too straightforward to be worth mapping any further. Other activities, such as 'rectify faulty equipment', may rely on the technician's skills and discretion to the extent that the activity has too much variation and is too complex to map in detail. Some activities, however, may need mapping in more detail to ensure quality or to protect the company's interests. For example, the activity of safety checking the customer's site to ensure that it is compliant with safety regulations will need specifying in some detail to ensure that the company can prove it exercised its legal responsibilities.

Process visibility

It is sometimes useful to map such processes in a way that makes the degree of visibility of each part of the process obvious. This allows those parts of the process with high visibility to be designed so that they enhance the customer's perception of the process. Figure 6.5 shows yet

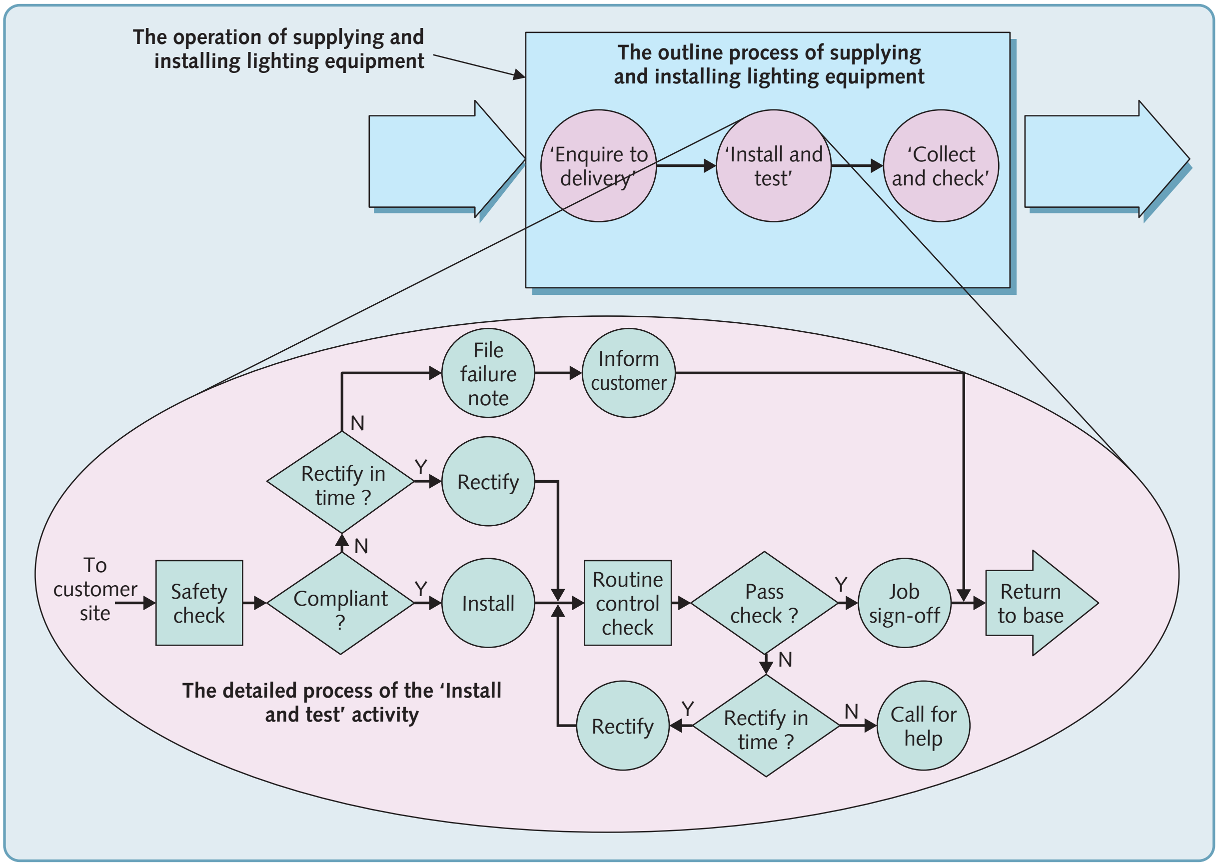

Figure 6.5 The 'supply and install' operations process mapped at three levels

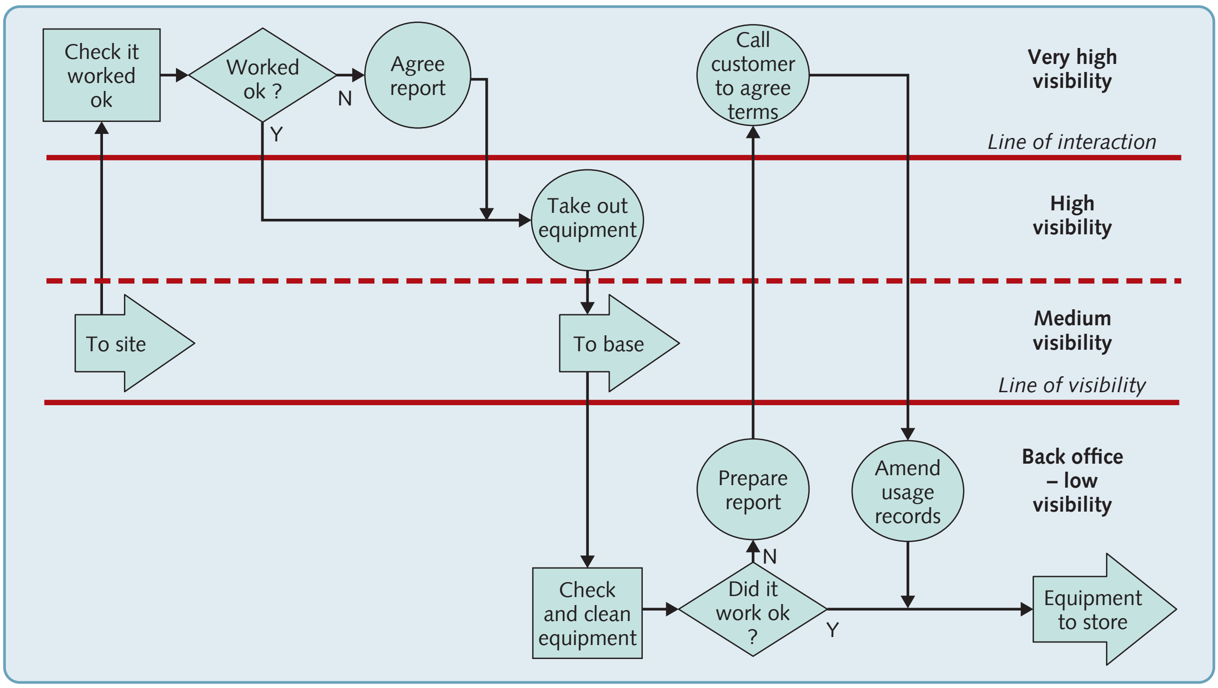

Figure 6.6 The 'collect and check' process mapped to show different levels of process visibility

another part of the lighting equipment company's operation: 'the collect and check' process. The process is mapped to show the visibility of each activity to the customer. Here four levels of visibility are used. There is no hard and fast rule about this; many processes simply distinguish between those activities that the customer *could* see and those that they couldn't. The boundary between these two categories is often called the 'line of visibility'. In Figure 6.6 three categories of visibility are shown. At the very highest level of visibility, above the 'line of interaction', are those activities that involve direct interaction between the lighting company's staff and the customer. Other activities take place at the customer's site or in the presence of the customer but involve less or no direct interaction. Yet further activities (the two transport activities in this case) have some degree of visibility because they take place away from the company's base and are visible to potential customers, but are not visible to the immediate customer.

EXAMPLE

Puncturing the line of visibility (by mistake)

Sometimes it gets embarrassing when customers see through the line of visibility. This happened when staff at Sainsbury's, a UK supermarket, mistakenly put up in its window a poster encouraging their workers to get customers to spend more. The poster, urging staff to get people to spend an extra 50p, appeared in a store in east London. It read: 'Fifty pence challenge – Let's encourage every customer to spend an additional 50p during each shopping trip between now and the year-end.' Unfortunately, before the mistake was noticed, a customer took a picture and posted it on Twitter saying: .@sainsburys not sure this is supposed to be in your window. Quickly Sainbury's tweeted back saying it should have remained behind closed doors and was meant for staff only. A spokesperson for Sainsbury's said: 'We often use posters to make store targets fun and achievable for our colleagues. They are intended for colleague areas in the store, but this one was mistakenly put on public display.'

Visibility, customer experience and emotional mapping

Processes with a high level of customer 'visibility' cannot be designed in the same way as processes that deal with inanimate materials or information. 'Processing' people is different. As we discussed in Chapter 1, operations and processes that primarily 'transform' people present a particular set of issues. Material and information are processed, but customers *experience* the process. When a customer experiences a process it results in them feeling emotions, not all of which are necessarily rational. Most of us have been made happy, angry, frustrated, surprised, reassured, or furious as customers in a process. Nor is the idea of considering how processes affect customer emotions confined to those processes that are intended to engage the emotions; for example as in entertainment-type organisations such as theme parks. Any high customer-contact product (or more likely, service) always creates an experience for the customer. Moreover, customer experience will affect customer satisfaction, and therefore has the potential to produce customer loyalty, influence expectations and create emotional bonds with customers. This is why many service organisations see how customers experience their processes (the so-called 'customer journey') at the core of their process design.

Designing processes with a significant experience content requires the systematic consideration of how customers may react to the experiences that the process exposes them to. This will include the sights, sounds, smells, atmosphere and general 'feeling' of the service. The concept of a 'servicescape', discussed in the previous chapter, is strongly related to consideration of engaging customers so that they connect with the process in a personal way. One of the most common methods of designing such processes is to consider what are commonly called 'touch-points'. These have been described as, 'everything the consumer uses to verify their service's effectiveness'.[5] They are the points of contact between a process and customers, and there might be many different touch-points during the customer journey. It is the accumulation of all the experiences from every touch-point interaction that shapes customers' judgement of the process.

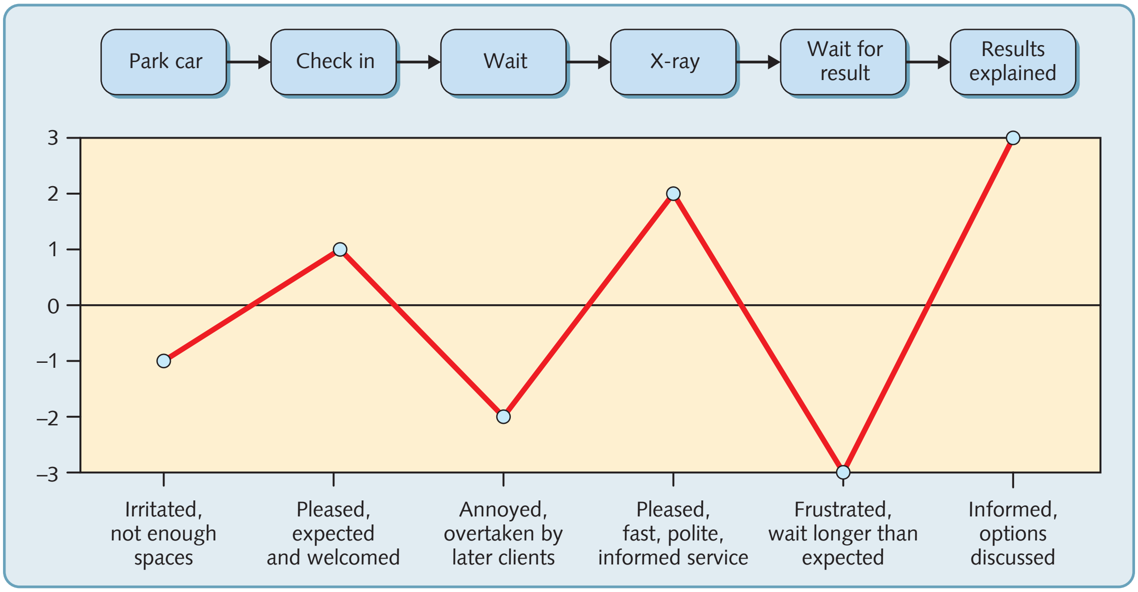

Figure 6.7 Customer experience map of a visit for an x-ray investigation

The features of a process at the touch-points are sometimes called 'clues' (or 'cues'); these are the messages that customers receive or experience as they progress through the process. The emotions that result from these cues contain the messages that the customer will receive and therefore influence how a customer will judge the process. When designing processes, managers need to ensure that all the messages coming from the clues at each stage of the process are consistent with the emotions they want the customers to experience and do not give them wrong or misleading messages about the process. In the same way as process mapping indicates the sequence and relationship between activities, so emotional mapping can indicate the type of emotions engendered in the customer's mind as they experienced the process. Figure 6.7 is a simplified version of how this might work for a visit to an x-ray clinic for a simple medical investigation. There are many ways that emotions can be mapped and different diagrammatic representations can be used. In this case a simple scoring system has been used ranging from +3 (very positive) to −3 (very negative). At each stage of the process, the reasons for the score are briefly noted.

OPERATIONS PRINCIPLE

The design of processes that deal with customers should consider the emotions engendered at each stage of the process.

DIAGNOSTIC QUESTION

Are process tasks and capacity configured appropriately?

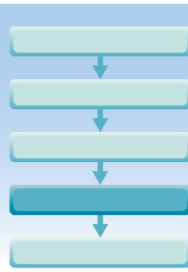

Process maps show how the activities of any particular process are currently arranged and help to suggest how they can be reconfigured. But there are also some general issues that must be understood before processes can be analysed. These relate to how the total task can be divided up within the process and determine how capacity is allocated. This, in turn, determines the flow through the process.

Getting to grips with process capacity means understanding the following issues:

- task precedence
- series and parallel configurations
- cycle time and process flow
- process balancing
- throughput, cycle time and work-in-process.

Task precedence

OPERATIONS PRINCIPLE
Process design must respect task precedence

Any process redesign needs to preserve the inherent precedence of activities within the overall task. Task 'precedence' defines what activities must occur before others, because of the nature of the task itself. At its simplest level, task precedence is defined by:

- the individual activities that comprise the total process task
- the relationship between these individual activities.

Task precedence is usually described by using a 'precedence diagram', which, in addition to the above, also includes the following information:

- the time necessary to perform the total task (sometimes known as the 'work content' of the task)
- the time necessary to perform each of the individual activities within the task.

EXAMPLE

Computer repair service centre

A repair service centre receives faulty or damaged computers sent in by customers, repairs them and dispatches them back to the customer. Each computer is subject to the same set of tests and repair activities, and although the time taken to repair each computer will depend on the results of the tests, there is relatively little variation between individual computers.

Table 6.3 defines the process task of testing and repairing the computers in terms of the seven activities that comprise the total task, the relationship between the activities in terms of each activity's 'immediate predecessor', and the time necessary to perform each activity. Figure 6.8 shows the relationship between the activities graphically. This kind of illustration is called the 'precedence diagram' for the process task. It is useful because it indicates how activities *cannot* be sequenced in the eventual process design. For example, the process cannot perform activity 'b' before activity 'a' is completed. However, it does not determine how a process *can* be designed. Yet once the task has been analysed in this way, activities can be arranged to form the processes' general configuration.

Table 6.3 Process task details for the 'computer test and repair' task

Activity code	*Activity name*	*Immediate predecessor*	*Activity time (mins)*
a	Preliminary test 1	-	5
c	Preliminary test 2	a	6
d	Dismantle	b	4
e	Test and repair 1	c	8
e	Test and repair 2	c	6
g	Test and repair 3	c	4
h	Clean/replace casing elements	d,e,f	10

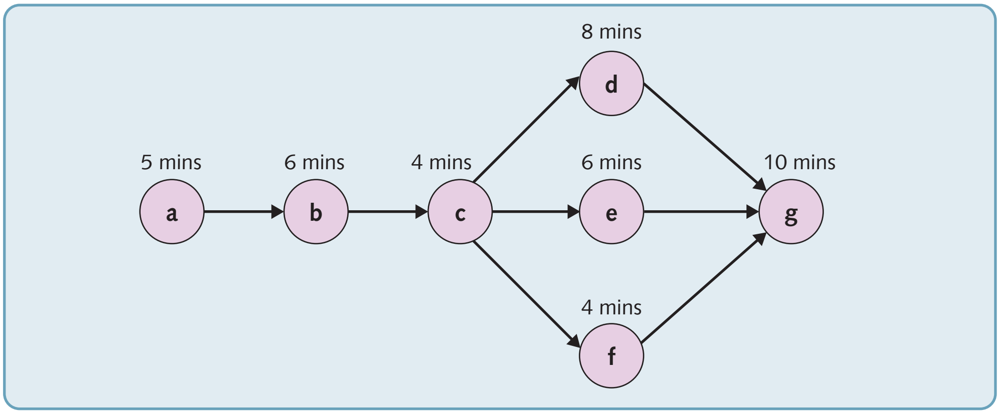

Figure 6.8 **Precedence diagram showing the relationship between activities for the computer test and repair task**

Series and parallel configurations

At its simplest level the general configuration of a process involves deciding the extent to which activities are arranged sequentially and the extent to which they are arranged in parallel.

For example, the task illustrated in Figure 6.8 involves seven activities that in total take 43 minutes. Demand is such that the process must be able to complete the test and repair task at the rate of one every 12 minutes in order to meet demand. One possible process design is to arrange the seven activities in a series arrangement of stages. The first question to address is, how many stages would this type of series arrangement require? This can be calculated by dividing the total work content of the task by the required cycle time.

In this case, number of stages = 43 minutes/12 minutes = 3.58 stages

Given the practical difficulties of having a fraction of a stage, this effectively means that the process needs four stages. The next issue is to allocate activities to each stage. Because the output from the whole process will be limited by the stage with most work (the sum of its allocated activities), each stage can have activities allocated to it up to a maximum allocation of 12 minutes. Figure 6.9 illustrates how this could be achieved. The longest stage (stage 2 in this case) will limit the output of the total process to one computer every 12 minutes and the other stages will be relatively under loaded.

However, there are other ways of allocating tasks to each stage, and involving the parallel arrangement of activities, that could achieve a similar output rate. For example, the four stages could be arranged as two parallel 'shorter' arrangements, with each stage performing approximately half of the activities in the total tasks. This is illustrated in Figure 6.10. It involves two, two-stage arrangements with stage 1 being allocated four activities that amount to 21 minutes of work and the second stage being allocated three activities that amount to 22 minutes of work. So, each arrangement will produce one repaired computer every 22 minutes (governed by the stage with the most work). This means that the two arrangements together will produce two repaired computers every 22 minutes, an average of one repaired computer every 11 minutes.

Loading each stage with more work and arranging the stages in parallel can be taken further. Figure 6.11 illustrates an arrangement where the whole test and repair task is performed at individual stages, all of which are arranged in parallel. Here, each stage will produce two repaired computer every 43 minutes and so together will produce four repaired computers every 43 minutes, an average output rate of one repaired computer every 10.75 minutes.

This simple example represents an important process design issue. Should activities in a process be arranged predominately in a single series 'long-thin' configuration, or, predominantly in several 'short-fat' parallel configurations, or somewhere in between? (Note that 'long' means the number of stages and 'fat' means the amount of work allocated to each stage.) Most

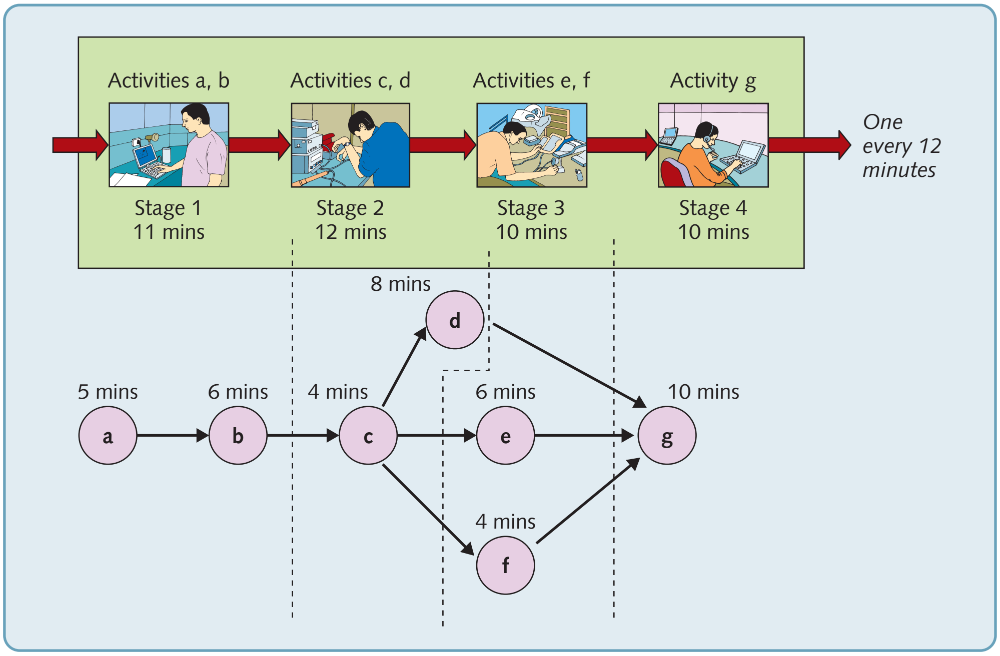

Figure 6.9 'Long-thin' arrangement of stages

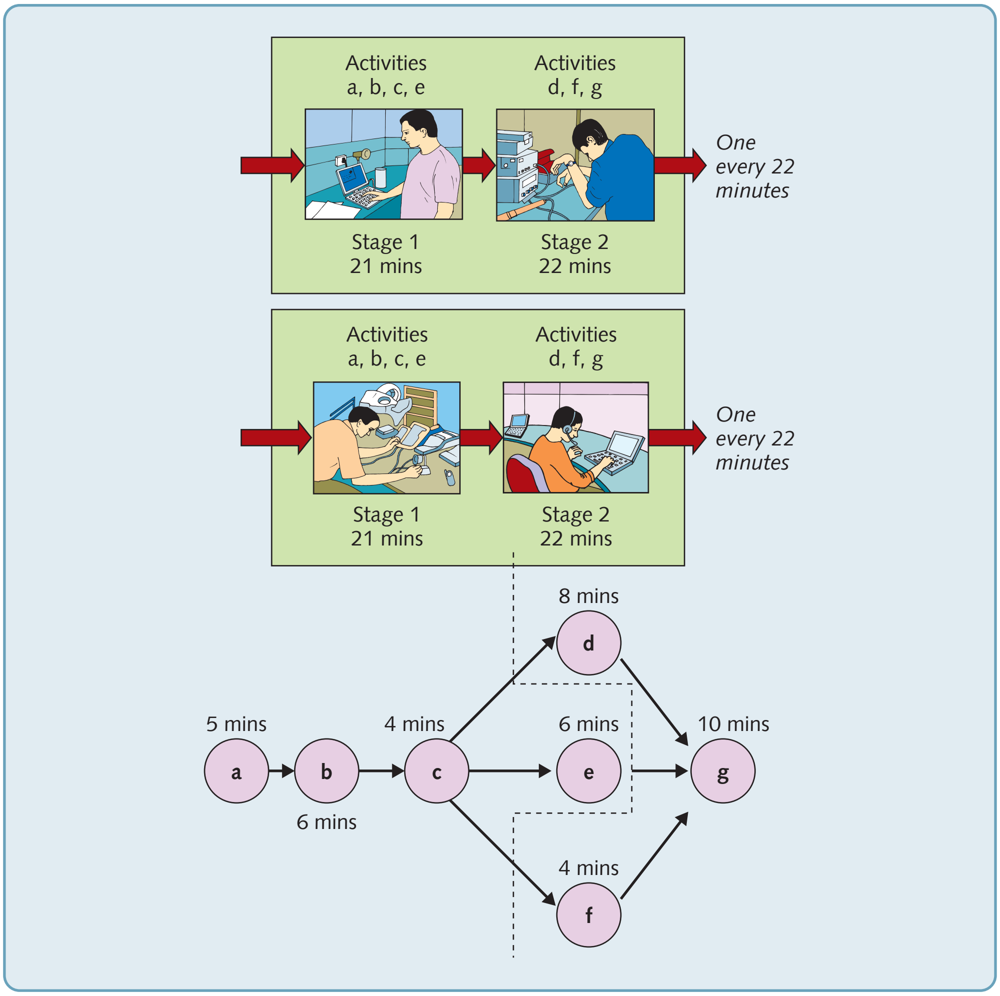

Figure 6.10 Intermediate configurations for the 'computer test and repair' task

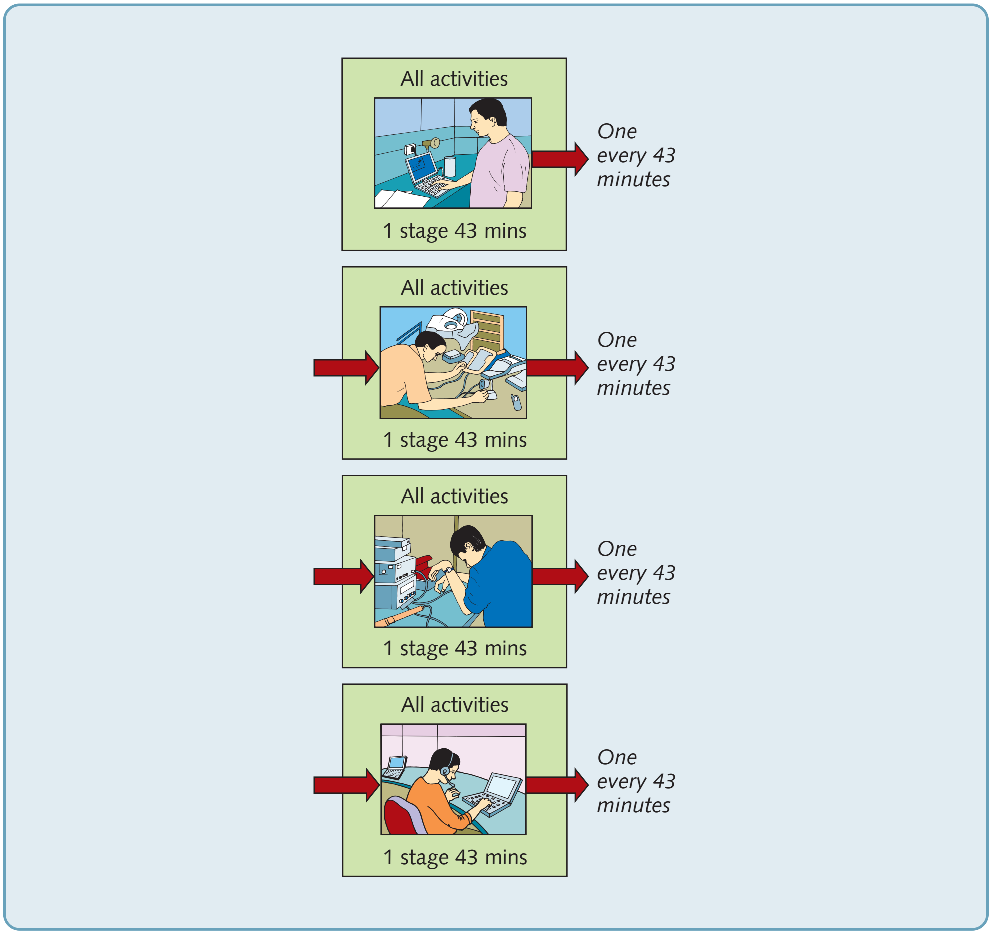

Figure 6.11 'Short-fat' configurations of stages

processes will adopt a combination of series and parallel configurations, and in any particular situation there are usually technical constraints which limit either how 'long and thin' or how 'short and fat' the process can be. But there is usually a real choice to be made with a range of possible options. The advantages of each extreme of the long-thin to short-fat spectrum are very different and help to explain why different arrangements are adopted.

The advantages of the series dominated (long-thin) configuration include:

- A more controlled flow through the process that is relatively easy to manage.
- Simple materials handling – especially if a product being manufactured is heavy, large or difficult to move.
- Lower capital requirements. If a specialist piece of equipment is needed for one element in the job, only one piece of equipment would need to be purchased; on short-fat arrangements every stage would need one.
- More efficient operation. If each stage is only performing a small part of the total job, the person at the stage may have a higher proportion of direct productive work as opposed to the non-productive parts of the job, such as picking up tools and materials.

The advantages of the parallel dominated (short-fat) configuration include:

- Higher mix flexibility. If the process needs to produce several types of product or service, each stage could specialise in different types.
- Higher volume flexibility. As volume varies, stages can simply be closed down or started up as required.

- Higher robustness. If one stage breaks down or ceases operation in some way, the other parallel stages are unaffected; a long-thin arrangement would cease operating completely.
- Less monotonous work. In the computer repair example, the staff in the short-fat arrangement repeat their tasks only every 43 minutes; in the long-thin arrangement it is every 12 minutes.

Cycle time and process capacity

The cycle time of a process is the time between completed units emerging from it. Cycle time is a vital factor in process design and has a significant influence on most of the other detailed design decisions. It is usually one of the first things to be calculated because it can be used both to represent the demand placed on a process and the processes' capacity. The cycle time also sets the pace or 'drum beat' of the process. However the process is designed it must be able to meet its required cycle time. It is calculated by considering the likely demand for the products or services over a period and the amount of production time available in that period.

OPERATIONS PRINCIPLE

Process analysis derives from an understanding of the required process cycle time.

EXAMPLE

Passport office

Suppose the regional government office that deals with passport applications is designing a process that will check applications and issue the documents. The number of applications to be processed is 1,600 per week and the time available to process the applications is 40 hours per week.

$$\text{Cycle time for the process} = \frac{\text{time available}}{\text{number to process}} = \frac{40}{1{,}600} = 0.025 \text{ hours} = 1.5 \text{ minutes}$$

So the process must be capable of dealing with a completed application once every 1.5 minutes, or 40 per hour.

Process capacity

If the cycle time indicates the output that must be achieved by a process, the next decision must be how much capacity is needed by the process in order to meet the cycle time. To calculate this, a further piece of information is needed – the work content of the process task. The larger the work content of the process task and the smaller the required cycle time, the more capacity will be necessary if the process is to meet the demand placed upon it.

EXAMPLE

Passport office

For the passport office, the total work content of all the activities that make up the total task of checking, processing and issuing a passport is, on average, 30 minutes.

So, a process with one person would produce a passport every 30 minutes. That is, one person would achieve a cycle time of 30 minutes. Two people would achieve a cycle time of $30/2 = 15$ minutes, and so on.

Therefore the general relationship between the number of people in the process (its capacity in this simple case) and the cycle time of the process is:

$$\frac{\text{Work content}}{N} = \text{cycle time}$$

Where: N = *the number of people in the process*

Therefore, in this case:

$$N = \frac{30}{\text{Cycle time}}$$

$$N = \frac{30}{1.5} = 20 \text{ people}$$

In this case:

So, the capacity that this process needs if it is to meet demand is 20 people.

Process balancing

Balancing a process involves attempting to allocate activities to each stage as equally as possible. Because the cycle time of the whole process is limited by the longest allocation of activity times to an individual stage, the more equally work is allocated, the less time will be 'wasted' at the other stages in the process. In practice it is nearly always impossible to achieve perfect balance so some degree of imbalance in work allocation between stages will occur. The effectiveness of the balancing activity is measured by balancing loss. This is the time wasted through the unequal allocation of activities as a percentage of the total time invested in processing. This is illustrated in Figure 6.12. Here the computer test and repair task is used to illustrate balancing loss for the 'long-thin' arrangement of four sequential stages and the 'intermediate' arrangement of two parallel two-stage arrangements.

OPERATIONS PRINCIPLE

Allocating work equally to each stage in a process (balancing) smooths flow and avoids bottlenecks.

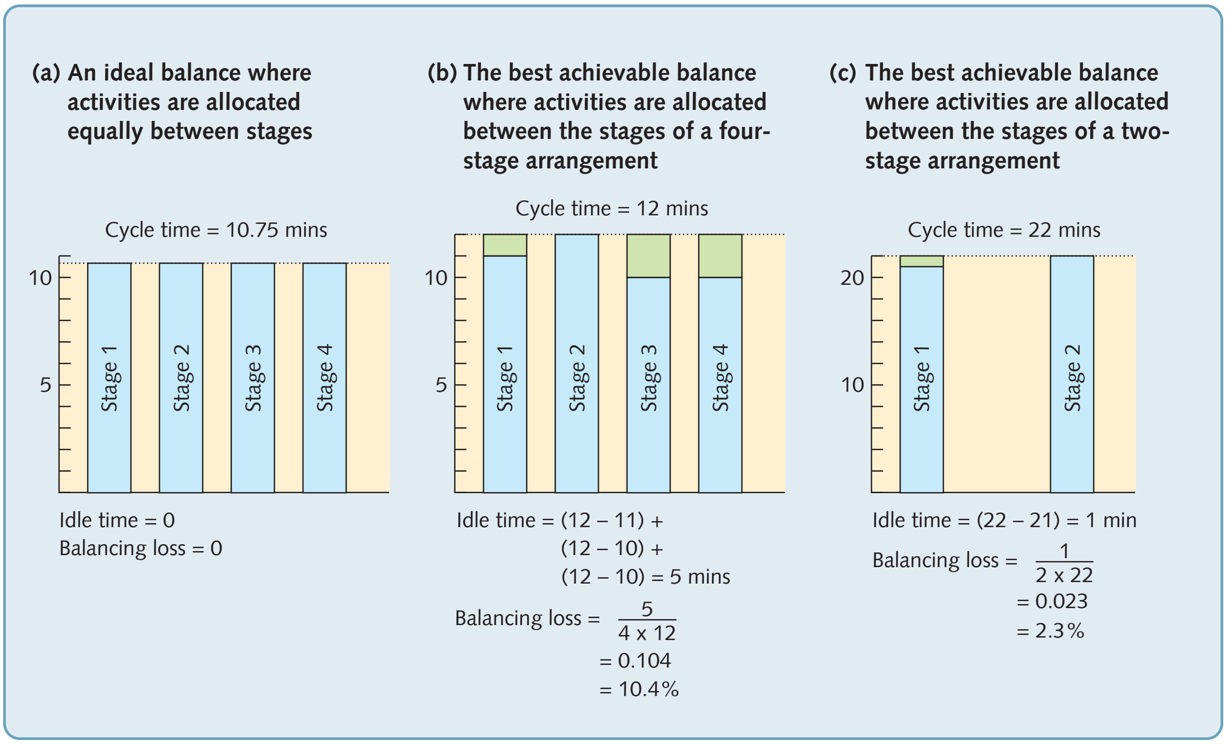

Figure 6.12 Balancing loss is that proportion of the time invested in processing the product or service that is not used productively

Figure 6.12(a) shows the ideal allocation of activities with each stage perfectly balanced. Here exactly a quarter of the total work content (10.75 minutes) has been allocated to each of the four stages. Every 10.75 minutes each stage performs its activities and passes a computer on to the next stage, or out of the process in the case of stage 4. No stage suffers any idle time and, because the stages are perfectly balanced, balancing loss = 0. In fact, because of the actual times of each activity, it is not possible to equally allocate work to each stage. Figure 6.12(b) shows the best allocation of activities. Most work is allocated to stage 2, so that stage will dictate the cycle time of the whole process. Stage one has only 11 minutes of work and so will be idle for (12 − 11) = 1 minute every cycle (or alternatively will keep processing one computer every 11 minutes and the build-up of inventory between stage 1 and stage 2 would grow to infinity). Similarly, stages three and four have idle time, in this case both have (12 − 10) = 2 minutes idle time. They can only process one computer every 12 minutes because stage 2 will only pass forward a computer to them every 12 minutes. So, they are being starved of work for 2 minutes every 12 minutes. In practice, stages that are not the bottleneck stage may not actually be idle for a period of time every cycle. Rather they will slow down the pace of work to match the time of the bottleneck stage. Nevertheless, this still is effective idle time because under conditions of perfect balance they could be performing useful work.

So, every cycle, all four stages are investing an amount of time equivalent to the cycle time to produce one completed computer. The total amount of invested time therefore is the number of stages in the process multiplied by the cycle time.

In this case, total invested time = 4 × 12 = 48 minutes

The total idle time for every computer processed is the sum of the idle times at the non-bottleneck stages, in this case 5 minutes.

Balancing loss is the amount of idle time as a percentage of the total invested time.

In this case, balancing loss = 5/(4 × 12) = 0.104 = 10.4%

Figure 6.12(c) makes the same calculation for the intermediate process described earlier. Here too, two-stage arrangements are placed in parallel. Stage 2 has the greatest allocation of work at 22 minutes, and will therefore be the bottleneck of the process. Stage 1 has 21 minutes' worth of work and therefore one minute of idle time every cycle.

Because the total invested time in the process each cycle = 2 × 22

The balancing loss = 1/(2 × 22) = 0.023 = 2.3%.

Throughput, cycle time and work-in-process

The cycle time of a process is a function of its capacity. For a given amount of work content in the process task, the greater the capacity of the process, the smaller its cycle time. In fact, the capacity of a process is often measured in terms of its cycle time, or more commonly the reciprocal of cycle time that is called 'throughput rate'. So, for example, A theme park ride would be described as having the capacity of 1,000 customers an hour, or an automated bottling line as having a capacity of 100 bottles a minute, and so on. However, a high level of capacity (short cycle time and fast throughput rate) does not necessarily mean that material, information or customers can move quickly through the process. This will depend on how many other units are contained within the process. If there is a large number of units within the process they may have to wait in 'work-in-process' inventories for part of the time they are within the process (throughput time).

Little's Law

The mathematical relationship that relates cycle time to work-in-process and throughput time is called Little's Law.[6] It is simple, but very useful, and it works for any stable process.

Little's Law can be stated as:

$$\text{Throughput time} = \text{Work-in-process} \times \text{Cycle time}$$

Or:

$$\text{Work-in-process} = \text{Throughput time} \times (1/\text{Cycle time})$$

That is:

$$\text{Work-in-process} = \text{Throughput time} \times \text{Throughput rate}$$

For example, in the case of the computer test and repair process with four stages:

Cycle time = 12 minutes (loading on the bottleneck station)

Work-in-process = 4 units (one at each stage of the process assuming there is no space for inventory to build up between stages)

OPERATIONS PRINCIPLE

Little's Law states that throughput time = work-in-process × cycle time.

Therefore, throughput time = work-in-process × cycle time

= 12 × 4 = 48 minutes

Similarly, for the example of the passport office, suppose the office has a 'clear desk' policy that means that all desks must be clear of work by the end of the day. How many applications should be loaded onto the process in the morning in order to ensure that every one is completed and desks are clear by the end of the day?

From before: cycle time = 1.5 minutes, and assuming 7.5 hour (450-minute) working day
From Little's Law:

throughput time = work-in-process × cycle time

450 minutes = work-in-process × 1.5

Therefore work-in-process = 450/1.5 = 300

So, 300 applications can be loaded onto the process in the morning and be cleared by the end of the working day.

EXAMPLE

Little's Law at a seminar

Mike was totally confident in his judgement, *'You'll never get them back in time,'* he said. *'They aren't just wasting time, the process won't allow them to all have their coffee and get back for 11 o'clock.'* Looking outside the lecture theatre, Mike and his colleague Dick were watching the 20 business men who were attending the seminar queuing to be served coffee and biscuits. The time was 10.45 and Dick knew that unless they were all back in the lecture theatre at 11 o'clock there was no hope of finishing his presentation before lunch. *'I'm not sure why you're so pessimistic,'* said Dick. *'They seem to be interested in what I have to say and I think they will want to get back to hear how operations management will change their lives.'* Mike shook his head, *'I'm not questioning their motivation,'* he said, *'I'm questioning the ability of the process out there to get through them all in time. I have been timing how long it takes to serve the coffee and biscuits. Each coffee is being made fresh and the time between the server asking each customer what they want and them walking away with their coffee and biscuits is taking 48 seconds. Remember that, according to Little's Law, throughput equals work-in-process multiplied by cycle time. If the work-in-process is the 20 managers in the queue and cycle time is 48 seconds, the total throughput time is going to be 20 multiplied by 0.8 minutes which equals 16 minutes. Add to that sufficient time for the last person to drink their coffee and you must expect a total throughput time of a bit over 20 minutes. You just haven't allowed long enough for the process.'* Dick was impressed. *'Err. . . what did you say that law was called again?' 'Little's Law,'* said Mike.

EXAMPLE

Little's Law at an IT support unit

Every year it was the same. All the workstations in the building had to be renovated (tested, new software installed, etc.) and there was only one week in which to do it. The one week fell in the middle of the August vacation period when the renovation process would cause minimum disruption to normal working. Last year the company's 500 workstations had all been renovated within one working week (40 hours). Each renovation last year took on average 2 hours and 25 technicians had completed the process within the week. This year there would be 530 workstations to renovate but the company's IT support unit had devised a faster testing and renovation routine that would only take on average 1.5 hours instead of 2 hours. How many technicians will be needed this year to complete the renovation processes within the week?

Last year:

$$\text{Work-in-process (WIP)} = 500 \text{ workstations}$$
$$\text{Time available } (T_t) = 40 \text{ hours}$$
$$\text{Average time to renovate} = 2 \text{ hours}$$
$$\text{Therefore, throughput rate } (T_r) = 1/2 \text{ per technician}$$
$$= 0.5N$$
$$\text{Where } N = \text{number of technicians}$$

Little's Law:

$$\text{WIP} = T_t \times T_r$$
$$500 = 40 \times 0.5\,N$$
$$N = \frac{500}{40 \times 0.5}$$
$$= 25 \text{ technicians}$$

This year:

$$\text{Work-in-process (WIP)} = 530 \text{ workstations}$$
$$\text{Time available} = 40 \text{ hours}$$
$$\text{Average time to renovate} = 1.5 \text{ hours}$$
$$\text{Throughput rate } (T_r) = 1/1.5 \text{ per technician}$$
$$= 0.67N$$
$$\text{where } N = \text{number of technicians}$$

Little's Law:

$$\text{WIP} = T_t \times T_r$$
$$530 = 40 \times 0.67N$$
$$N = \frac{530}{40 \times 0.67}$$
$$= 19.88 \text{ technicians}$$

DIAGNOSTIC QUESTION

Is process variability recognised?

So far in our treatment of process analysis we have assumed that there is no significant variability either in the demand to which the process is expected to respond, or in the time taken for the process to perform its various activities. Clearly, this is not the case in reality. So, it is important to look at the variability that can affect processes and take account of it. However, do not dismiss the deterministic analysis we have been examining up to this point. At worst it provides a good first approximation to analysing processes, while at best, the relationships that we have discussed do hold for average performance values.

Sources of variability in processes

There are many reasons why variability occurs in processes. A few of these possible sources of variation are listed below:

- The late (or early) arrival of material, information or customers at a stage within the process.
- The temporary malfunction or breakdown of process technology within a stage of the process.
- The necessity for recycling 'misprocessed' materials, information or customers to an earlier stage in the process.
- The misrouting of material, information or customers within the process that then needs to be redirected.
- Each product or service being processed might be different, for example, different models of automobile going down the same line.
- Products or services, although essentially the same, might require slightly different treatment. For instance, in the computer test and repair process, the time of some activities will vary depending on the results of the diagnostic checks.
- With any human activity there are slight variations in the physical coordination and effort on the part of the person performing the task that result in variation in activity times, even of routine activities.

All these sources of variation within a process will interact with each other, but result in two fundamental types of variability:

- Variability in the demand for processing at an individual stage within the process, usually expressed in terms of variation in the inter-arrival times of units to be processed.
- Variation in the time taken to perform the activities (i.e. process a unit) at each stage.

Activity time variability

The effects of variability within a process will depend on whether the movements of units between stages, and hence the inter-arrival times of units at stages, are synchronised or not. For example, consider the computer test and repair process described previously. Figure 6.13 shows the average activity time at each stage of the process, but also the variability around

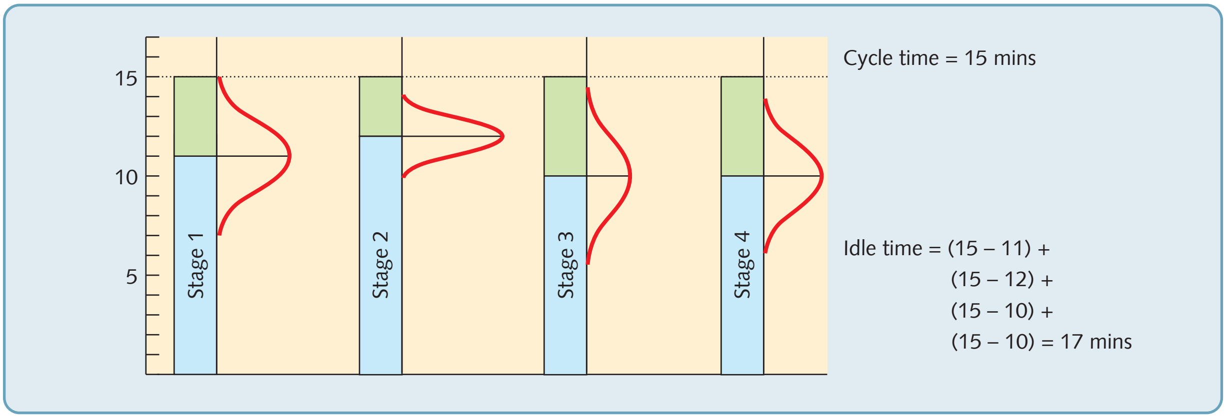

Figure 6.13 Processing time variability in a synchronised process. Cycle time will need to accommodate the longest activity time at any of the stages.

the average time. Suppose that it was decided to synchronise the flow between the four stages by using an indexing conveyor or a simple traffic lights system that ensured all movement between the stages happened simultaneously. The interval between each synchronised movement would have to be set at an interval that would allow all stages to have finished their activities irrespective of whether they had experienced a particularly fast or particularly slow activity time. In this case, from Figure 6.13 that synchronised indexing time would have to be set at 15 minutes. This then becomes the effective cycle time of the process. Note that the effective bottleneck stage is now stage 1 rather than stage 2. Although stage 2 has the longer average activity time (12 minutes), activity one with an average activity time of 11 minutes has a degree of variability that results in a maximum activity of 15 minutes. Note also that every stage will experience some degree of idle time, the average idle time at each station being the cycle time minus the average activity time at that station. This reduction in the efficiency of the process is only partly a result of its imbalance. The extra lost time is as a result of activity time variability.

OPERATIONS PRINCIPLE

Variability in a process acts to reduce its efficiency.

This type of effect is not at all uncommon. For example, automobiles are assembled using a moving belt assembly line whose speed is set to achieve a cycle time that can accommodate activity time variability. However, a more common arrangement, especially when processing information or customers, is to move units between stages as soon as the activities performed by each stage are complete. Here, units move through the process in an unsynchronised manner rather than having to wait for an imposed movement time. This means that each stage may spend less time waiting to move their unit forward, but it does introduce more variation in the demand placed on subsequent stations. When movement was synchronised the inter-arrival time of units at each stage was fixed at the cycle time. Without synchronisation, the inter-arrival time at each stage will itself be variable.

Arrival time variability

To understand the effect of arrival variability on process performance it is first useful to examine what happens to process performance in a very simple process as arrival time changes under conditions of no variability. For example, the simple process shown in Figure 6.14 comprises one stage that performs exactly 10 minutes of work. Units arrive at the process at a constant and predictable rate. If the arrival rate is one unit every 30 minutes, then the

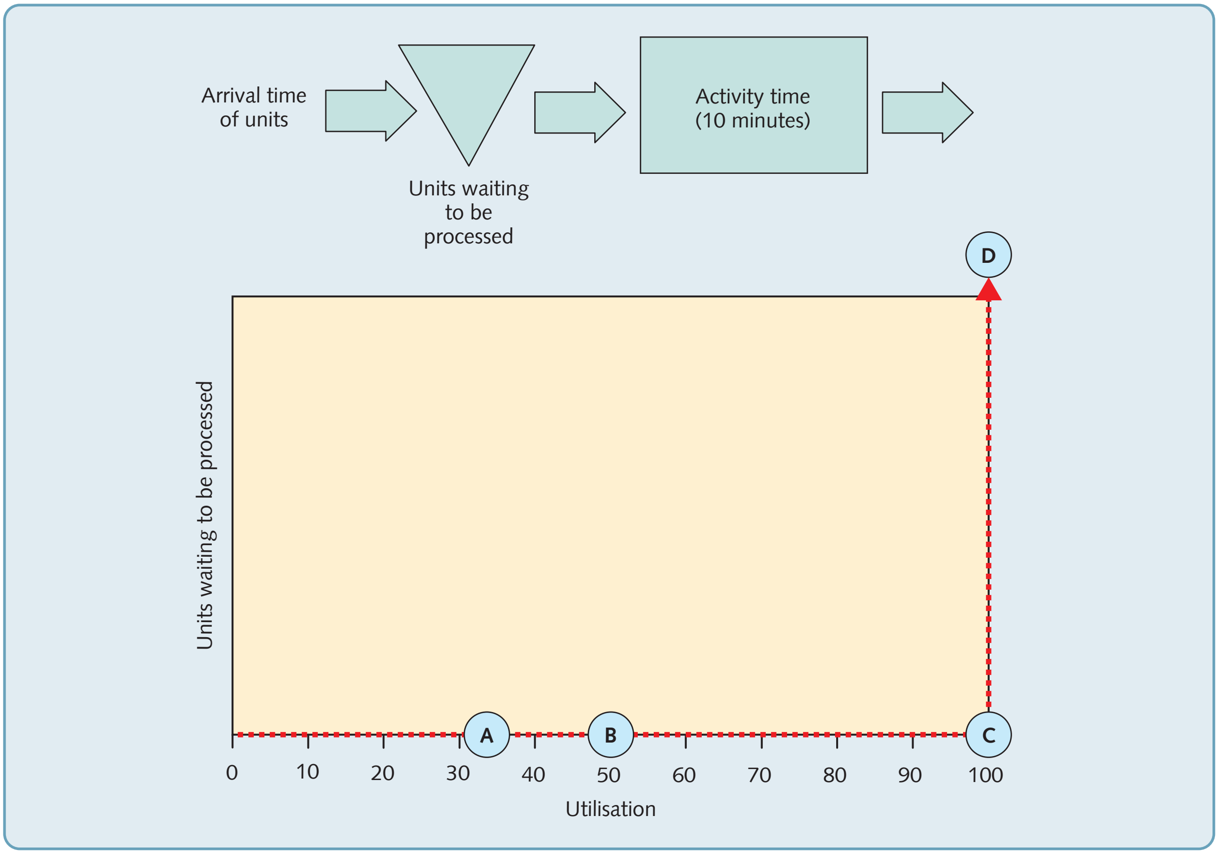

Figure 6.14 The relationship between process utilisation and number of units waiting to be processed for no arrival time or activity time variability

process will be utilised for only 33.33 per cent of the time, and the units will never have to wait to be processed. This is shown as point A on Figure 6.14. If the arrival rate increases to one arrival every 20 minutes, the utilisation increases to 50 per cent, and again the units will not have to wait to be processed. This is point B on Figure 6.14. If the arrival rate increases to one arrival every 10 minutes, the process is now fully utilised, but, because a unit arrives just as the previous one has finished being processed, no unit has to wait. This is point C on Figure 6.14. However, if the arrival rate ever exceeded one unit every 10 minutes, the waiting line in front of the process activity would build up indefinitely, as is shown as point D in Figure 6.14. So, in a perfectly constant and predictable world, the relationship between process waiting time and utilisation is a rectangular function as shown by the dotted line in Figure 6.14.

When arrival time is not constant but variable, then the process may have both units waiting to be processed and under-utilisation of the processes' resources over the same period. Figure 6.15 illustrates how this could happen for the same process as shown in Figure 6.14, with constant 10-minute activity times, but this time with variable arrival times. The table gives details of the arrival of each unit at the process and when it was processed, and the bar chart illustrates it graphically. Six units arrive with an average arrival time of 11 minutes, some of which can be processed as soon as they arrive (units A, D and F) while others have to wait for a short period. Over the same period the process has to wait for work three times.

OPERATIONS PRINCIPLE

Process variability results in simultaneous waiting and resource under utilisation.

Unit	Arrival time	Start of activity	End of activity	Wait time
A	0	0	10	0
B	12	12	22	0
C	20	22	32	2
D	34	34	44	0
E	43	44	54	1
F	55	55	65	0

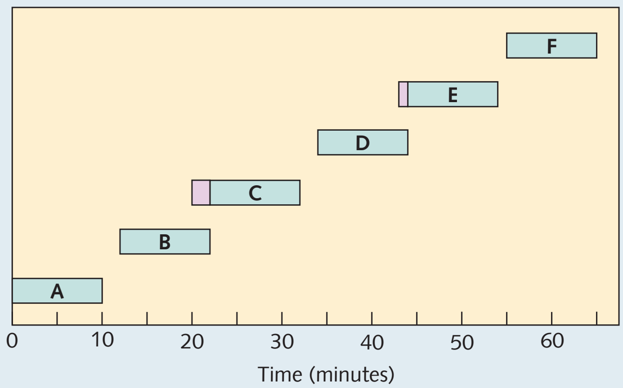

Figure 6.15 Units arriving at a process with variable arrival times and a constant activity time (10 minutes)

During the observed period:

Time when a single unit was waiting = 3 minutes

Elapsed time for processing the six units = 65 minutes

Average number of units waiting = 3/65

= 0.046 units

Process idle time = 5 minutes

So, process idle percentage = 5 × 100/65

= 7.7%

Therefore, process utilisation = 92.3%

This point is shown as point X in Figure 6.16. If the average arrival time were to be changed with the same variability, the dotted line in Figure 6.16 would show the relationship between average waiting time and process utilisation. As the process moves closer to 100 per cent utilisation the higher the average waiting time will become. Or, to put it another way, the only way to guarantee very low waiting times for the units is to suffer low process utilisation.

When both arrival times and activity times are variable, this effect is even more pronounced. And the greater the variability, the more the waiting time utilisation deviates from the simple rectangular function of the 'no variability' conditions that was shown in Figure 6.14. A set of curves for a typical process is shown in Figure 6.17(a). This phenomenon has important implications for the design of processes. In effect it presents three options to process designers wishing to improve the waiting time or utilisation performance of their processes, as shown in Figure 6.17(b). Either:

- accept long average waiting times and achieve high utilisation (point X)
- accept low utilisation and achieve short average waiting times (point Y)
- reduce the variability in arrival times, activity times, or both, and achieve higher utilisation and short waiting times (point Z).

To analyse processes with both inter-arrival and activity time variability queuing or 'waiting line' analysis can be used. This is treated in the supplement to this chapter. But, do not dismiss the relationship shown in Figures 6.16 and 6.17 as some minor technical phenomenon. It is far more than this. It identifies an important choice in process design that could have strategic

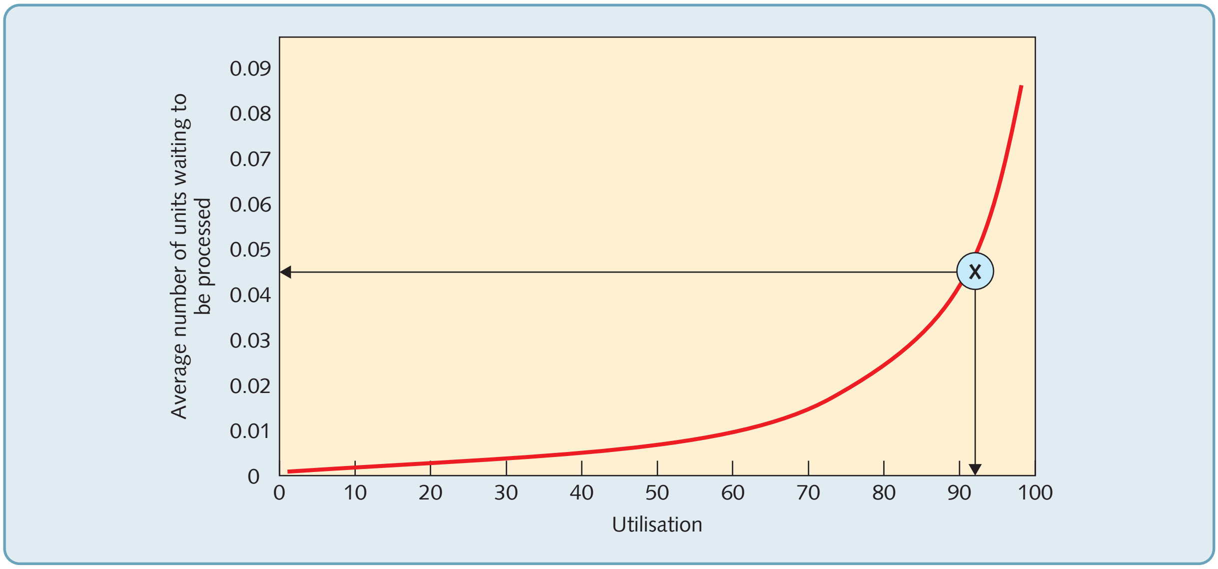

Figure 6.16 The relationship between process utilisation and number of units waiting to be processed for the variable arrival times in the example

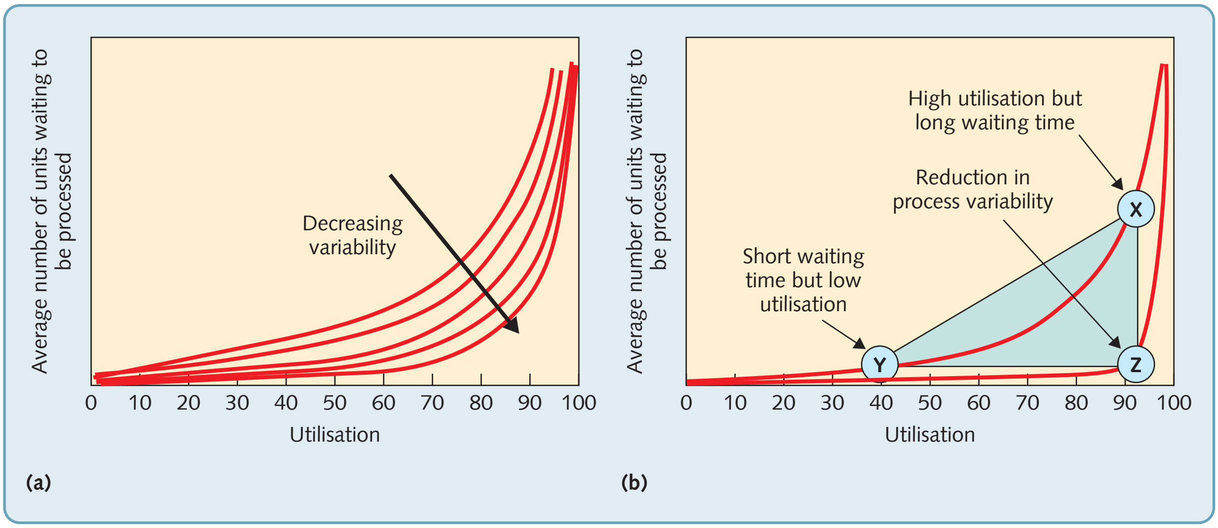

Figure 6.17 The relationship between process utilisation and number of units waiting to be processed for variable arrival and activity times.

implications. Which is more important to a business, fast throughput time, or high utilisation of its resources? The only way to have both of these simultaneously is to reduce variability in its processes, which may itself require strategic decisions such as limiting the degree of customisation of products or services, or imposing stricter limits on how products or services can be delivered to customers, and so on. It also demonstrates an important point concerned with the day-to-day management of process – the only way to absolutely guarantee 100 per cent utilisation of resources, is to accept an infinite amount of work-in-process and/or waiting time. We will take this point further in Chapter 8 when we deal with capacity management.

OPERATIONS PRINCIPLE

Process design involves some choice between utilisation, waiting time and variability reduction.

Critical commentary

Unless you are one of those that instinctively reacts against the very idea of 'processes', there is not too much that would be considered contentious in this chapter. However, some practitioners would reject the idea of mapping processes as they exist currently. Rather, they would advocate a more radical 'clean sheet of paper' approach. Only by doing this, they would say, could one be sufficiently imaginative in the redesign of processes. Having said that, the details of process analysis are not too contentious; there is one criticism that challenges the whole basis of how we think about processes. All the processes used as illustrations in this chapter involve relatively simple stage-by-stage arrangements very much like an 'assembly line' where work is linearly synchronised to deliver a completed task. Yet many business processes are far more complex than this with fuzzy and simultaneous activities that interact in unpredictable ways. And, although we have made the point that not all processes should be designed in the same way, there is relatively little guidance in conventional process design theory about how to tackle these more complex processes.

SUMMARY CHECKLIST

- ☐ Have a clear set of performance objectives for each process been set?
- ☐ Do the process design objectives clearly relate to the business's strategic objectives?
- ☐ Is the following information known for all key processes in the operation?
  - The throughput or flow rate of the process?
  - The throughput time of the process?
  - The number of units in the process (work-in-process)?
  - The utilisation of process resources?
- ☐ Are processes documented using process mapping techniques?
- ☐ Are the emotional responses of any customers / people who are being 'processed' understood?
- ☐ Are formal process descriptions followed in practice?
- ☐ If not, should the process descriptions be changed or should existing process descriptions be enforced?
- ☐ Is it necessary for process descriptions to include the degree of visibility at each stage of the process?
- ☐ Are the details of task precedence known for each process?
- ☐ Have the advantages and disadvantages of series and parallel configurations been explored?
- ☐ Is the process balanced? If not, can the bottleneck stages be redesigned to achieve better balance?
- ☐ Are the relationships between throughput, cycle time and work-in-process understood (Little's Law)?
- ☐ Are the sources of process variability recognised?
- ☐ Has the effect of variability been recognised in the design of the process?

CASE STUDY

The Action Response Applications Processing Unit (ARAPU)

Introduction

Action Response is a London-based charity dedicated to providing fast responses to critical situations throughout the world. It was founded by Susan N'tini, its Chief Executive, to provide relatively short-term aid for small projects until they could obtain funding from larger donors. The charity receives requests for cash aid usually from an intermediary charity and looks to process the request quickly, providing funds where and when they are needed. *'Give a man a fish and you feed him today, teach him to fish and you feed him for life, it's an old saying and it makes sense but, and this is where Action Response comes in, he might starve while he's training to catch fish.'*(Susan N'tini)

Nevertheless, Susan does have some worries. She faces two issues in particular. First she is receiving complaints that funds are not getting through quickly enough. Second the costs of running the operation are starting to spiral. She explains. '*We are becoming a victim of our own success. We have striven to provide greater accessibility to our funds; people can access application forms via the internet, by post and by phone. But we are in danger of losing what we stand for. It is taking longer to get the money to where it is needed and our costs are going up. We are in danger of failing on one of our key objectives: to minimise the proportion of our turnover that is spent on administration. At the same time we always need to be aware of the risk of bad publicity through making the wrong decisions. If we don't check applications thoroughly, funds may go to the "wrong" place and if the newspapers gets hold of the story we would run a real risk of losing the goodwill, and therefore the funds, from our many supporters*'.

Susan held regular meetings with key stakeholders. One charity that handled a large number of applications for people in Nigeria told her of frequent complaints about the delays over the processing of the applications. A second charity representative complained that when he telephoned to find out the status of an application the ARAPU staff did not seem to know where it was or how long it might be before it was complete. Furthermore, he felt that this lack of information was eroding his relationship with his own clients some of whom were losing faith in him as a result, '*trust is so important in the relationship*', he explained.

Some of Susan's colleagues, while broadly agreeing with her anxieties over the organisation's responsiveness and efficiency, took a slightly different perspective. *'One of the really good things about Action Response is that we are more flexible than most charities. If there a need and if they need support until one of the larger charities can step in, then we will always consider a request for aid. I would not like to see any move towards high process efficiency harming our ability to be open-minded and consider requests that might seem a little unusual at first.'* (Jacqueline Horton, Applications Assessor)

Others saw the charity as performing an important counselling role. *'Remember that we have gained a lot of experience in this kind of short-term aid. We are often the first people that are in a position to advise on how to apply for larger and longer term funding. If we developed this aspect of our work we would again be fulfilling a need that is not adequately supplied at the moment.'* (Stephen Nyquist, Applications Assessor)

The Action Response Applications Processing Unit (ARAPU)

Potential aid recipients, or the intermediary charities representing them, apply for funds using a standard form. These forms can be downloaded from the internet or requested via a special help line. Sometimes the application will come directly from an individual community leader but more usually it will come via an intermediary charity that can help the applicant to complete the form. The application is sent to ARAPU, usually by fax or post (some were submitted online, but few communities have this facility).

ARAPU employs seven applications assessors with support staff who are responsible for data entry, coding, filing and 'completing' (staff who prepare payment, or explain why no aid can be given). In addition, a board of non-paid trustees meets every Thursday, to approve the assessors' decisions. The unit's IT system maintained records of all transactions, providing an update on the number of applications received, approved, declined, and payments allocated. These reports identified that the unit received about 300 new applications per week and responded to about the same number (the unit operates a 35 hour week). But while the unit's financial targets were being met, the trend indicated that cost per application was increasing. The target for the turnaround of an application, from receipt of application to response was 20 days, and although this was not measured formally, it was generally assumed that turnaround time was longer than this. Accuracy had never been an issue, as all files were thoroughly assessed to ensure that all the relevant data were collected before the applications were processed. Productivity seemed high and there was always plenty of work waiting for processing at each section, with the exception that the 'completers' were sometimes waiting for work to come from the committee on a Thursday. Susan had conducted an inspection of all sections' in-trays that had revealed a rather shocking total of about 2,000 files waiting within the process, not counting those waiting for further information.

Processing applications

The processing of applications is a lengthy procedure requiring careful examination by applications assessors, trained to make well-founded assessments in line with the charity's guidelines and values. Incoming application are opened by one of the four 'receipt' clerks who check that all the necessary forms have been included in the application; the receipt clerks take about 10 minutes per application. These are then sent to the coding staff, in batches twice a day. The five coding clerks allocate a unique identifier to each application and key the information on the application into the system. The coding stage takes about 20 minutes for each application. Files are then sent to the senior applications assessors' secretary's desk. As assessors become available, the secretary provides the next job in the line to the assessor.

About one hundred of the cases seen by the assessors each week are put aside after only 10 minutes 'scanning' because the information is ambiguous, so further information is needed. The assessor returns these files to the secretaries, who write to the applicant (usually via the intermediate charity) requesting additional information, and return the file to the 'receipt' clerks who 'store' the file until the further information eventually arrives (usually between 1 and 8 weeks). When it does arrive, the file enters the process and progresses through the same stages again. Of the applications that require no further information, around half (150) are accepted and half (150) declined. On average, those applications that were not 'recycled' took around 60 minutes to assess.

All the applications, whether approved or declined, are stored prior to ratification. Every Thursday the Committee of Trustees meets to formally approve the applications assessors' decisions. The committee's role is to sample the decisions to ensure that the guidelines of the charity are upheld. In addition, they will review any particularly unusual cases highlighted by the applications assessors. Once approved by the committee the files are then taken to the completion officers. There are three 'decline' officers whose main responsibility is to compile a suitable response to the applicant, pointing out why the application failed and offering, if possible, helpful advice. An experienced declines officer takes about 30 minutes to finalise the file and write a suitable letter. Successful files are passed to the four 'payment' officers where again the file is completed, letters (mainly standard letters) are created and payment instructions are given to the bank. This usually takes around 50 minutes, including dealing with any queries from the bank about payment details. Finally, the paperwork itself is sent, with the rest of the file, to two 'dispatch' clerks who complete the documents and mail them to the applicant. The dispatch activity takes, on average, 10 minutes for each application.

The feeling amongst the staff was generally good. When Susan consulted the team they said their work was clear and routine, but their life was made difficult by charities that rang in expecting them to be able to tell them the status of an application they had submitted. It could take them hours, sometimes days, to find any individual file. Indeed, two of the 'receipt' clerks now were working almost full time on this activity. They also said that charities frequently complained that decision-making seemed slow.

QUESTIONS

1 What objectives should the ARAPU process be trying to achieve?

2 What is the main problem with the current ARAPU processes?'

3 How could the ARAPU process be improved?

APPLYING THE PRINCIPLES

Some of these exercises can be answered by reading the chapter. Others will require some general knowledge of business activity and some might require an element of investigation. **Model answers for the first two questions are available on the student companion website.** For model answers to the other questions in this section, please ask your tutor/lecturer.

1 One of the examples at the beginning of the chapter described 'drive-through' fast food processes. Think about (or better still, visit) a drive-though service and try mapping what you can see of the process (plus what you can infer from what may be happening 'behind the scenes').

2 *'It is a real problem for us'*, said Angnyeta Larson, *'We now have only ten working days between all the expense claims coming from the departmental coordinators and authorizing payments on the next month's payroll. This really is not long enough and we are already having problems during peak times.'* Angnyeta was the department head of the internal financial control department of a metropolitan authority in southern Sweden. Part of her department's responsibilities included checking and processing expense claims from staff throughout the metropolitan authority and authorising payment to the salaries payroll section. She had 12 staff who were trained to check expense claims and all of them were devoted full time to processing the claims in the two weeks (10 working days) prior to the deadline for informing the salaries section. The number of claims submitted over the year averaged around 3,200, but this could vary between 1,000 during the quiet summer months up to 4,300 in peak months. Processing claims involved checking receipts, checking that claims met with the strict financial allowances for different types of expenditure, checking all calculations, obtaining more data from the claimant if necessary, and (eventually) sending an approval notification to salaries. The total processing time took on average 20 minutes per claim.

(a) How many staff does the process need on average, for the lowest demand, and for the highest demand?

(b) If a more automated process involving electronic submission of claims could reduce the average processing time to 15 minutes, what effect would this have on the required staffing levels?

(c) If department coordinators could be persuaded to submit their batched claims earlier (not always possible for all departments) so that the average time between submission of the claims to the finance department and the deadline for informing salaries section was increased to 15 working days, what effect would this have?

3 The headquarters of a major creative agency offered a service to all its global subsidiaries that included the preparation of a budget estimate that was submitted to potential clients when making a 'pitch' for new work. This service had been offered previously only to a few of the groups subsidiary companies. Now that it was to be offered worldwide, it was deemed appropriate to organise the process of compiling budget estimates on a more systematic basis. It was estimated that the worldwide demand for this service would be around 20 budget estimates per week, and that, on average, the staff who would put together these estimates would be working a 35-hour week. The elements within the total task of compiling a budget estimate are shown in Table 6.4.

(a) What is the required cycle time for this process?

(b) How many people will the process require to meet the anticipated demand of 20 estimates per week?

(c) Assuming that the process is to be designed on a 'long-thin' basis, what elements would each stage be responsible for completing? And what would be the balancing loss for this process?

(d) Assuming that instead of the long-thin design, two parallel processes are to be designed, each with half the number of stations of the long-thin design, what now would be the balancing loss?

Table 6.4 The elements within the total task of compiling a budget estimate

Element	*Time (mins)*	*What element (s) must be done prior to this one?*
A – obtain time estimate from creatives	20	None
B – obtain account handler's deadlines	15	None
C – obtain production artwork estimate	80	None
D – preliminary budget calculations	65	A, B and C
E – check on client budget	20	D
F – check on resource availability and adjust estimate	80	D
G – complete final budget estimate	80	E and F

4 A company has decided to manufacture a general-purpose 'smoothing plane', a tool which smooths and shapes wood. Its engineers estimated the time it would take to perform each element in the assembly process. The marketing department also estimated that the likely demand for the new product would be 98,000 units. The marketing department was not totally confident of its forecast; however, *'a substantial proportion of demand is likely to be export sales, which we find difficult to predict. But whatever demand does turn out to be, we will have to react quickly to meet it. The more we enter these parts of the market, the more we are into impulse buying and the more sales we lose if we don't supply."*

An idea of the assembly task can be gained from Table 6.5 that gives the 'standard time' for each element of the assembly task

Table 6.5 Standard times for each element of assembly task in standard minutes (SM)

Press elements	
Assemble poke	0.12 min
Fit poke to front	0.10 min
Rivet adjusting lever to front	0.15 min
Press adjusting nut screw to front	0.08 min
Bench elements	
Fit adjusting nut to front	0.15 min
Fit frog screw to front	0.05 min
Fit knob to base	0.15 min
Fit handle to base	0.17 min
Fit front assembly to base	0.15 min
Assemble blade unit	0.08 min
Final assembly	0.20 min
Packing element	
Make box, wrap plane, pack	0.20 min
Total time	**1.60 min**

All elements must be performed sequentially in the order listed.

The standard costing system at the company involves adding a 150 per cent overhead charge to the direct labour cost of manufacturing the product, and the product would retail for the equivalent of around €35 in Europe where most retailers will sell this type of product for about 70–100 per cent more than they buy it from the manufacturer.

(a) How many people will be needed to assemble this product?

(b) Design a process for the assembly operation (to include the fly press work) including the tasks to be performed at each part of the system.

(c) How might the process design need to be adjusted as demand for this and similar products builds-up?

5 At the theatre, the interval during a performance of 'King Lear' lasts for 20 minutes and in that time 86 people need to use the toilet cubicles. On average, a person spends 3 minutes in the cubicle. There are 10 cubicles available. (a) Does the theatre have enough toilets to deal with the demand? (b) If there are not enough cubicles, how long should the interval be to cope with demand?

6 A gourmet burger shop has a daily demand for 250 burgers and operates for 10 hours. (a) What is the required cycle time in minutes? (b) Assuming that each burger has 7.2 minutes of work required, how many servers are required? (c) If the burger shop has a three-stage process for making burgers. Stage 1 takes 2.0 minutes, stage 2 takes 3.0 minutes, and stage 3 takes 2.2 minutes, what is the balancing loss for the process?

Notes on chapter

1 Sources include: Zhang, S. (2016) 'How to fit the world's biggest indoor waterfall in an airport', *Wired,* 9 July ; *Airport Technology* (2014) 'Terminal 4, Changi International Airport, Singapore', airporttechnology.com; SAA Architects (2013) 'South East Asia Aviation Hub Development - design led Changi Airport Terminal 4, Singapore', saa.com, 6 November; Driver, C. (2014) 'And the winners are . . . Singapore crowned the best airport in the world (and Heathrow scoops top terminal)', Mail online, 28 March.

2 Sources include: Inc.well blog (2016) 'The differences you'll notice the next time you order at a McDonald's drive-thru', nbcchicago.com, 23 November; Oches, S. (2013) 'The drive-thru performance study', *QSA Magazine,* October; Horovitz, A. (2002) 'Fast food world ways drive-thru is the way to go', *USA Today,* 3 April.

3 Sources include: Cornwall, S. (2013) 'Ecover announces world-first in plastic packaging', *Packaging Gazette,* 7 March; Ecover website, http://www.the-splash.co.uk/; Osborne, H. (2006). 'Spick'n'span ethics', *The Guardian* 17 November, http://www.guardian.co.uk/environment/2006/nov/17/

4 Not everyone uses exactly the same terminology in this area. For example, some publications use the term 'cycle time' to refer to what we have called 'throughput rate'.

5 Shostack, L. G. (1984) 'Designing services that deliver', *Harvard Business Review,* 62(1), 133–139.

6 Little's law is best explained in Hopp, W. J. and Spearman, M. L. (2001) *Factory Physics* (2nd edn), McGrawHill.

7 Earlier editions of this book (mistakenly) used North American notation that defines cycle time and processing rates slightly differently. The notation used here is consistent with that used earlier in this chapter.

TAKING IT FURTHER

Damelio, R. (2011) *The Basics of Process Mapping* (*2nd edn*), *Productivity Press.* *A practitioner book that is both very comprehensive and up-to-date.*

Hammer, M. (1990) *'Reengineering Work: Don't automate, obliterate', Harvard Business Review, July-August.* *This is the paper that launched the whole idea of business processes and process management in general to a wider managerial audience. Slightly dated but worth reading.*

Harrington, H.J. (2011) *Streamlined Process Improvement*, *McGraw Hill Professional.* *Practical and insightful.*

Harvard Business Review (2011) *'Improving Business Processes' (Harvard Pocket Mentor) Harvard Business School Press.* *A collection of HBR papers.*

Hopp, W.J. and Spearman, M.L. (2011) *Factory Physics* (*3rd edn*), *Waveland Pr Inc.* *Very technical so don't bother with it if you aren't prepared to get into the maths. However, some fascinating analysis, especially concerning Little's Law.*

Ramaswamy, R. (1996) *Design and Management of Service Processes, Addison-Wesley Longman.* *A relatively technical approach to process design in a service environment.*

Slack, N. (2017) *The Operations Advantage, Kogan Page.* *The chapter on 'internal processes' expands on some of the issues discussed here.*

Sparks, W. (2016) *Process Mapping Road Trip: Improve organizational workflow in five steps, Promptitude Publishing.* *A practitioners guide – straightforward and sensible.*

SUPPLEMENT TO CHAPTER 6

Queuing analysis

Queuing analysis (in many parts of the world it is called 'waiting line' analysis) is often explained purely in terms of customers being processed through service operations. This is misleading. Although queuing analysis can be particularly important in service operations, especially where customers really do 'queue' for service, the approach is useful in any kind of operation. Figure 6.18 shows the general form of queuing analysis. Customers arrive according to some probability distribution and wait to be processed (unless part of the operation is idle); when they have reached the front of the queue, they are processed by one of the *n* parallel 'servers' (their processing time also being described by a probability distribution), after which they leave the operation. There are many examples of this kind of system. Table 6.6 illustrates some of these. All of these examples can be described by a common set of elements that define their queuing behaviour:

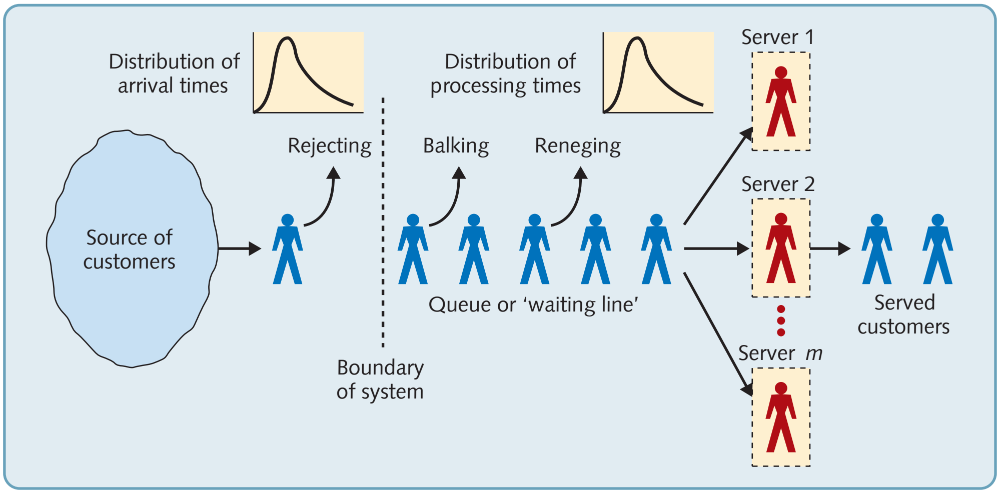

Figure 6.18 The general form of queuing analysis

Table 6.6 Examples of processes that can be analysed using queuing analysis

Operation	*Arrivals*	*Processing capacity*
Bank	Customers	Tellers
Supermarket	Shoppers	Checkouts
Hospital clinic	Patients	Doctors
Graphic artist	Commissions	Artists
Custom cake decorators	Orders	Cake decorators
Ambulance service	Emergencies	Ambulances with crews
Telephone switchboard	Calls	Telephonists
Maintenance department	Breakdowns	Maintenance staff

The source of customers, sometimes called the calling population is the source of supply of customers. In queue management 'customers' are not always human. 'Customers' could for example be trucks arriving at a weighbridge, orders arriving to be processed or machines waiting to be serviced, etc.

The arrival rate is the rate at which customers needing to be served arrive at the server or servers. Rarely do customers arrive at a steady and predictable rate. Usually, there is variability in their arrival rate. Because of this it is necessary to describe arrival rates in terms of probability distributions.

The queue: customers waiting to be served form the queue or waiting line itself. If there is relatively little limit on how many customers can queue at any time, we can assume that, for all practical purposes, an infinite queue is possible. Sometimes, however, there is a limit to how many customers can be in the queue at any one time.

Queue discipline: this is the set of rules that determine the order in which customers waiting in the queue are served. Most simple queues, such as those in a shop, use a *first-come-first-served* queue discipline.

Servers: a server is the facility that processes the customers in the queue. In any queuing system there may be any number of servers configured in different ways. In Figure 6.14 servers are configured in parallel, but some systems may have servers in a series arrangement. There is also likely to be variation in how long it takes to process each customer. Therefore processing time, like arrival time, is usually described by a probability distribution.

Calculating queue behaviour

Management scientists have developed formulae that can predict the steady-state behaviour of different types of queuing system. Unfortunately, many of these formulae are extremely complicated, especially for complex queuing systems, and are beyond the scope of this book. In fact, in practice, computer programs are almost always used to predict the behaviour of queuing systems. However, studying queuing formulae can illustrate some useful characteristics of the way queuing systems behave. Moreover, for relatively simple systems, using the formulae (even with some simplifying assumptions) can provide a useful approximation to process performance.

Notation

There are several different conventions for the notation used for different aspects of queuing system behaviour. It is always advisable to check the notation used by different authors before using their formulae. We shall use the following notation.

t_a = average time between arrival

r_a = arrival rate (items per unit time) = $1/t_a$

c_a = coefficient of variation of arrival times

m = number of parallel servers at a station

t_e = mean processing time

r_e = processing rate (items per unit time) = m/t_e

c_e = coefficient of variation of process time

u = utilisation of station = $r_a/r_e = (r_a t_e)/m$

WIP = average work-in-process (number of items) in the queue

WIP = expected work-in-process (number of times) in the queue

W_q = expected waiting time in the queue

W = *expected waiting time in the system* (queue time + processing time)

Variability

The concept of variability is central to understanding the behaviour of queues. If there were no variability there would be no need for queues to occur because the capacity of a process could be relatively easily adjusted to match demand. For example, suppose one member of staff (a server) serves customers at a bank counter who always arrive exactly every 5 minutes (i.e. 12 per hour). Also suppose that every customer takes exactly 5 minutes to be served, then because,

(a) the arrival rate is ≤ processing rate, and
(b) there is no variation

No customer need ever wait because the next customer will arrive when, or before, the previous customer. That is, $WIP_q = 0$.

Also, in this case, the server is working all the time, again because exactly as one customer leaves, the next one is arriving. That is, $u = 1$.

Even with more than one server, the same may apply. For example, if the arrival time at the counter is 5 minutes (12 per hour) and the processing time for each customer is now always exactly 10 minutes, the counter would need two servers, and because,

(a) arrival rate is ≤ processing rate m, and
(b) there is no variation

Again, $WIP_q = 0$, and $u = 1$.

Of course, it is convenient (but unusual) if arrival rate/processing rate = a whole number. When this is not the case (for this simple example with no variation)

$$\text{Utilisation} = \text{processing rate}/(\text{arrival rate} \times m)$$

For example, if arrival rate, r_a = 5 minutes
processing rate, r_e = 8 minutes
number of servers, m = 2
then, utilisation, $u = 8/(5 \times 2) = 0.8$ or 80%

Incorporating variability

The previous examples were not realistic because they assumed no variation in arrival or processing times. We also need to take into account the variation around these means. To do that, we need to use a probability distribution. Figure 6.19 contrasts two processes with different arrival distributions. The units arriving are shown as people, but they could be jobs arriving at a machine, trucks needing servicing or any other uncertain event. The top example shows low variation in arrival time where customers arrive in a relatively predictable manner. The bottom

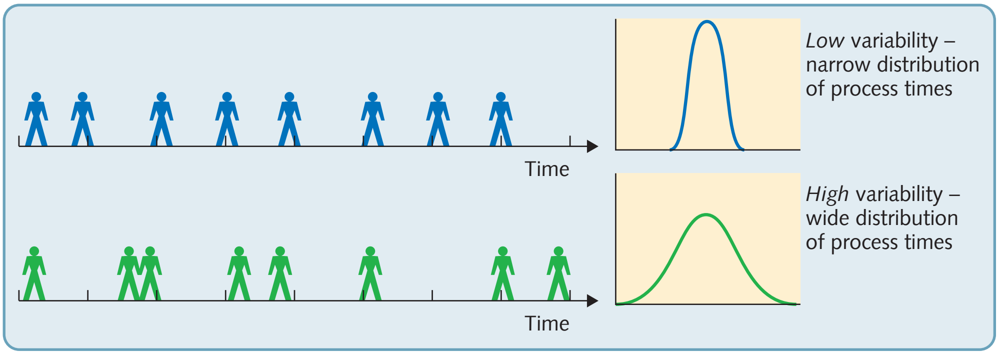

Figure 6.19 Low and high arrival variation

example has the same average number of customer arriving but this time they arrive unpredictably with sometimes long gaps between arrivals and at other times two or three customers arriving close together. We could do a similar analysis to describe processing times.

In Figure 6.19 high arrival variation has a distribution with a wider spread (called 'dispersion') than the distribution describing lower variability. Statistically the usual measure for indicating the spread of a distribution is its standard deviation, σ. But variation does not only depend on standard deviation. For example, a distribution of arrival times may have a standard deviation of 2 minutes. This could indicate very little variation when the average arrival time is 60 minutes. But it would mean a very high degree of variation when the average arrival time is 3 minutes. Therefore to normalise standard deviation, it is divided by the mean of its distribution. This measure is called the coefficient of variation of the distribution. So,

$$c_a = \text{coefficient of variation of arrival times} = \sigma_a / t_a$$

$$c_e = \text{coefficient of variation of processing times} = \sigma_e / t_e$$

Incorporating Little's Law

Little's Law (described earlier in the chapter) describes the relationship between the cycle time, the work-in-process and the throughput time of the process. It was denoted by the following simple relationship.[7]

$$\text{Throughput time} = \text{Work-in-process} \times \text{cycle time}$$

So

$$\text{Working in process} = \text{Throughput time/cycle time}$$

Or

$$\text{WIP} = T/C$$

And

$$\text{Cycle time} = 1/\text{arrival rate}$$

Little's Law can help to understand queuing behaviour. Consider the queue in front of station.

Work-in-process in the queue = arrival rate at the queue (equivalent to 1/cycle time) × waiting time in the queue (equivalent to throughput time)

$$\text{WIP}_q = r_a \times W_q$$

and

waiting time in the whole system = the waiting time in the queue + the average process time at the station

$$W = W_q + t_e$$

We will use this relationship later to investigate queuing behaviour.

Types of queuing system

Conventionally queuing systems are characterised by four parameters.

A = the distribution of arrival times (or more properly inter-arrival times, the elapsed times between arrivals)
B = the distribution of process times
m = the number of servers at each station
b = the maximum number of items allowed in the system

The most common distributions used to describe A or B are either,

(a) the exponential (or Markovian) distribution denoted by M
(b) the general (for example normal) distribution denoted by G.

So, for example, an M / G / 1 /5 queuing system would indicate a system with exponentially distributed arrivals, process times described by a general distribution such as a normal distribution, with one server and a maximum number of five items allowed in the system. This type of notation is called Kendall's Notation.

Queuing theory can help us investigate any type of queuing system, but in order to simplify the mathematics, we shall here deal only with the two most common situations. Namely:

M/M/m - the exponential arrival and processing times with *m* servers and no maximum limit to the queue.
G/G/m - general arrival and processing distributions with *m* servers and no limit to the queue.

And first we will start by looking at the simple case when M = 1.

For M/M/1 queuing systems

The formulae for this type of system are as follows:

$$\text{WIP} = \frac{u}{1-u}$$

Using Little's Law,

$$\text{WIP} = \text{Throughput time/cycle time}$$

$$\text{Throughput time} = \text{WIP} \times \text{cycle time} = \text{WIP/arrival rate}$$

Then:

$$\text{Throughput time} = \frac{u}{1-u} \times \frac{1}{r_a} = \frac{t_e}{1-u}$$

and since, throughput time in the queue = total throughput time − average processing time:

$$\begin{aligned} W_q &= W - t_e \\ &= \frac{t_e}{1-u} - t_e \\ &= \frac{t_e - t_e(1-u)}{1-u} = \frac{t_e - t_e - ut_e}{1-u} \\ &= \frac{u}{(1-u)} t_e \end{aligned}$$

again, using Little's Law:

$$\mathrm{WIP}_q = r_a \times W_q = \frac{u}{(1-u)} t_e\, r_a$$

and since:

$$u = \frac{r_a}{r_e} = r_a\, t_e$$

$$r_a = \frac{u}{t_e}$$

then:

$$\begin{aligned} \mathrm{WIP}_q &= \frac{u}{(1-u)} \times t_e \times \frac{u}{t_e} \\ &= \frac{u^2}{(1-u)} \end{aligned}$$

For M/M/m systems

When there are m servers at a station the formula for waiting time in the queue (and therefore all other formulae) needs to be modified. Again, we will not derive these formulae but just state them.

$$W_q = \frac{u^{\sqrt{2(m+1)}-1}}{m(1-u)} t_e$$

From which the other formulae can be derived as before.

For G/G/1 systems

The assumption of exponential arrival and processing times is convenient as far as the mathematical derivations of various formulae are concerned. However, in practice, process times in particular are rarely truly exponential. This is why it is important to have some idea of how a G/G/1 and G/G/M queue behaves. However, exact mathematical relationships are not possible with such distributions. Therefore some kind of approximation is needed. The one here is in common use, and although it is not always accurate, it is for practical purposes. For G/G/1 systems the formula for waiting time in the queue is as follows.

$$W_q = \left(\frac{c_a^2 + c_e^2}{2}\right)\left(\frac{u}{(1-u)}\right) t_e$$

There are two points to make about this equation. The first is that it is exactly the same as the equivalent equation for an M/M/1 system but with a factor to take account of the variability of the arrival and process times. The second is that this formula is sometimes known as the VUT formula because it describes the waiting time in a queue as a function of,

V = the variability in the queuing system
U = the utilisation of the queuing system (that is demand versus capacity) and
T = the processing times at the station.

In other words, we can reach the intuitive conclusion that queuing time will increase as variability, utilisation or processing time increase.

For G/G/m systems

The same modification applies to queuing systems using general equations and m servers. The formula for waiting time in the queue is now as follows.

$$W_q = \left(\frac{c_a^2 + c_e^2}{2}\right)\left(\frac{u^{\sqrt{2(m+1)}-1}}{m(1-u)}\right)t_e$$

WORKED EXAMPLE

'I can't understand it. We have worked out our capacity figures and I am sure that one member of staff should be able to cope with the demand. We know that customers arrive at a rate of around 6 per hour and we also know that any trained member of staff can process them at a rate of 8 per hour. So why is the queue so large and the wait so long? Have at look at what is going on there please.'

Sarah knew that it was probably the variation, both in customers arriving and in how long it took each of them to be processed, that was causing the problem. Over a two-day period when she was told that demand was more or less normal, she timed the exact arrival times and processing times of every customer. Her results were as follows:

The coefficient of variation, c_a of customer arrivals = 1
The coefficient of variation, c_e of processing time = 3.5
The average arrival rate of customers, r_a = 6 per hour
Therefore the average inter-arrival time = 10 minutes
The average processing rate, r_e = 8 per hour
Therefore the average processing time = 7.5 minutes
Therefore the utilisation of the single server, u = 6/8 = 0.75

Using the waiting time formula for a G/G/1 queuing system

$$W_q = \left(\frac{1 + 12.25}{2}\right)\left(\frac{0.75}{1 - 0.75}\right)7.5$$

$$= 6.625 \times 3 \times 7.5 = 149.06 \text{ mins}$$

$$= 2.48 \text{ hours}$$

Also because:

$$\text{WIP}_q = \text{cycle time} \times \text{throughput time}$$

$$\text{WIP}_q = 6 \times 2.48 = 14.68$$

So, Sarah had found out that the average wait that customers could expect was 2.48 hours and that there would be an average of 14.68 people in the queue.

'Ok, so I see that it's the very high variation in the processing time that is causing the queue to build up. How about investing in a new computer system that would standardise processing time to a greater degree? I have been talking with our technical people and they reckon that,

if we invested in a new system, we could cut the coefficient of variation of processing time down to 1.5. What kind of a difference would this make?'

Under these conditions with $c_e = 1.5$

$$W_q = \left(\frac{1 + 12.25}{2}\right)\left(\frac{0.75}{1 - 0.75}\right) 7.5$$

$$= 1.625 \times 3 \times 7.5 = 36.56 \text{ mins}$$

$$= 0.61 \text{ hours}$$

Therefore:

$$\text{WIP}_q = 6 \times 0.61 = 3.66$$

In other words, reducing the variation of the process time has reduced average queuing time from 2.48 hours down to 0.61 hours and has reduced the expected number of people in the queue from 14.68 down to 3.66.

WORKED EXAMPLE

A bank wishes to decide how many staff to schedule during its lunch period. During this period, customers arrive at a rate of 9 per hour and the enquiries that customers have (such as opening new accounts, arranging loans etc.) take on average 15 minutes to deal with. The bank manager feels that four staff should be on duty during this period but wants to make sure that the customers do not wait more than 3 minutes on average before they are served. The manager has been told by his small daughter that the distributions that describe both arrival and processing times are likely to be exponential. Therefore,:

$$r_a = 9 \text{ per hour, therefore}$$

$$t_a = 6.67 \text{ minutes}$$

$$r_e = 4 \text{ per hour, therefore}$$

$$t_e = 15 \text{ minutes}$$

The proposed number of servers, $m = 4$ therefore the utilisation of the system, $u = 9/4 \times 4 = 0.5625$.

From the formula for waiting time for a M/M/m system,

$$W_q = \frac{u^{\sqrt{2(m+1)}-1}}{m(1 - u)} t_e$$

$$W_q = \frac{u^{\sqrt{2(m+1)}-1}}{m(1 - u)} t_e$$

$$W_q = \frac{0.5625^{\sqrt{10}-1}}{4(1 - 0.5625)} \times 0.25$$

$$= \frac{0.5625^{2.162} - 1}{1.75} \times 0.25$$

$$= 0.042 \text{ hours}$$

$$= 2.52 \text{ minutes}$$

$$W_q = \frac{0.5625^{\sqrt{10}-1}}{4(1 - 0.5625)} \times$$

$$= \frac{0.5625^{2.162} - 1}{1.75} \times 0.25$$

$$= 0.042 \text{ hours}$$

$$= 2.52 \text{ minutes}$$

Therefore the average waiting time with four servers would be 2.52 minutes, that is well within the manager's acceptable waiting tolerance.

TAKING IT FURTHER

Hopp, W.J. and Spearman, M.L. (2001) *Factory Physics (2nd edn)* , *McGraw Hill.* *Very technical so don't bother with it if you aren't prepared to get into the maths. However, some fascinating analysis, especially concerning Little's Law.*

7 Supply chain management

Introduction

Supply chain management involves all the things that organisations do to get their products and services to their customers. It is important because it directly affects their ability to deliver products or services into the market. It is why businesses such as Apple, Toyota and Zara, achieve notable results in competitive markets. Chapter 4 treated the strategic design of supply networks. This chapter considers the planning and control activity for the individual supply chains in the network. Supply chain management is the overarching operations management activity that dictates an operation's *delivery* performance because it controls the flow of products and services from suppliers right through to the end customer. That is why it is the first chapter dealing with the planning and control of delivery. But planning and controlling delivery is a much larger topic, and includes capacity management (Chapter 8), inventory management (Chapter 9), resource planning and control (Chapter 10) and lean synchronization (Chapter 11). Figure 7.1 illustrates the supply–demand linkage treated in this chapter.

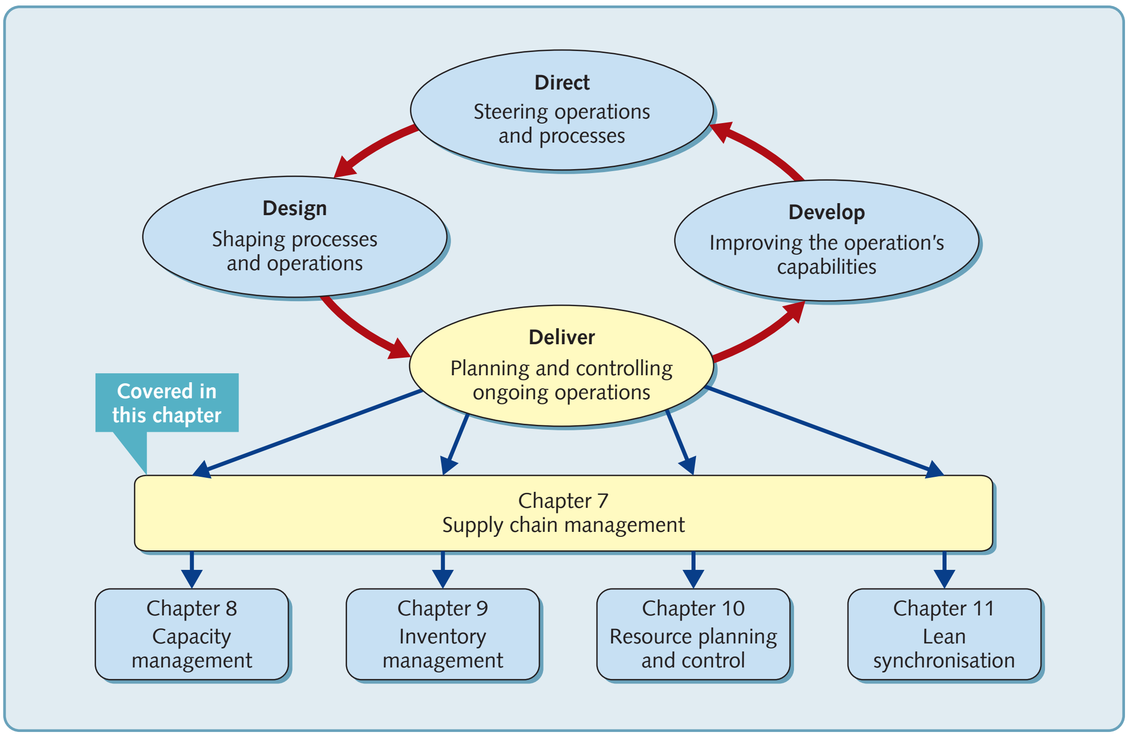

Figure 7.1 Supply chain management is the management of the relationships and flows between operations and processes; it is the topic that integrates all the issues concerning the delivery of products and services

EXECUTIVE SUMMARY

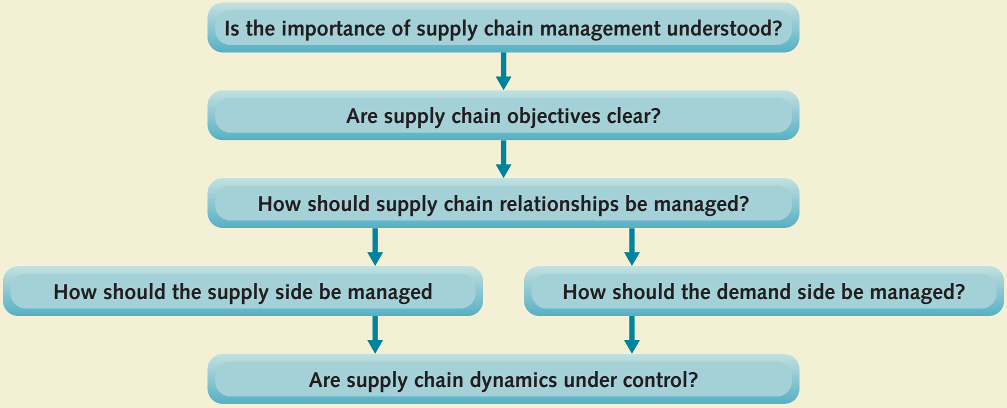

Is the importance of supply chain management understood?

The end consequence of what customers experience when they purchase a product or service is the result, not of an individual operation, but of a whole supply chain of operations. So the reputation and commercial success of an individual company depends fundamentally on its position in, and relationship with, the other operations in its supply chain. No operations manager can afford to ignore it. Technically, supply network management looks at all the operations or processes in a network, while supply chain management refers to a string of operations or processes. However, the two terms are often used interchangeably. Many of the principles of managing external supply chains (flow between operations) are also applicable to internal supply chains (flow between processes).

Are supply chain objectives clear?

The central objective of supply chain management is to satisfy the needs of the end customer. So, each operation in the chain should contribute to whatever mix of quality, speed, dependability, flexibility and cost that the end customer requires. Individual operations failure in any of these objectives can be multiplied throughout the chain. So, although each operation's performance may be adequate, the performance of the whole chain could be poor. An important distinction is between lean and agile supply chain performance. Broadly, lean (or efficient) supply chains are appropriate for stable 'functional' products and services, while agile (or responsive) supply chains are more appropriate for less predictable innovative products and services.

How should supply chain relationships be managed?

Supply chain relationships can be described on a spectrum from market-based, contractual, 'arm's-length' relationships, through to close and long-term partnership relationships. Each has its advantages and disadvantages. Developing relationships involves assessing which relationship will provide the best potential for developing overall performance. However, the types of relationships adopted may be dictated by the structure of the market itself. If the number of potential suppliers is small, there are few opportunities to use market mechanisms to gain any kind of advantage.

How is the supply side managed?

Managing supply side relationships involves three main activities: selecting appropriate suppliers; planning and controlling ongoing supply activity; and supplier development. Supplier selection involves trading off different supplier attributes, often using scoring assessment methods. Managing ongoing supply involves clarifying supply expectations, often using service-level agreements to manage the supply relationships. Supplier development can benefit both suppliers and customers, especially in partnership relationships. Very often barriers are the mismatches in perception between customers and suppliers.

How is the demand side managed?

This will depend partly on whether demand is dependent on some known factor and therefore predictable, or independent of any known factor and therefore less predictable. Approaches such as materials requirements planning (MRP) are used in the former case, while approaches such as inventory management are used in the latter case. The increasing outsourcing of physical distribution and the use of new tracking technologies, such as radio frequency identification (RFID), have brought efficiencies to the movement of physical goods and customer service. But customer service may be improved even more if suppliers take on responsibility for customer development; that is, helping customers to help themselves.

Are supply chain dynamics under control?

Supply chains have a dynamic of their own that is often called the *bullwhip* effect. It means that relatively small changes at the demand end of the chain increasingly amplify into large disturbances as they move upstream. Three methods can be used to reduce this effect. Information sharing can prevent overreaction to immediate stimuli and give a better view of the whole chain. Channel alignment through standardised planning and control methods allows for easier coordination of the whole chain. Improving the operational efficiency of each part of the chain prevents local errors multiplying to affect the whole chain.

DIAGNOSTIC QUESTION

Is the importance of supply chain management understood?

Supply chain management (SCM) is the management of the relationships and flows between the 'string' of operations and processes that produce value in the form of products and services to the ultimate consumer. It is a holistic approach to managing across the boundaries of companies and of processes. Technically, supply *chains* are different to supply *networks*. A supply network is *all* the operations that are linked together so as to provide goods and services through to end customers. In large supply networks there can be many hundreds of supply chains of linked operations passing through a single operation. The same distinction holds within operations. Internal supply network, and supply chain, management concerns flow between processes or departments. See Figure 7.2. Confusingly, the terms supply network and supply chain management are often (mistakenly) used interchangeably. It is also worth noting that the 'flows' in supply chains are not restricted to the downstream flow of products and services from suppliers through to customers. Although the most obvious failure in supply chain management occurs when downstream flow fails to meet customer requirements, the root cause may be a failure in the upstream flow of information. Modern supply chain management is as much concerned with managing information flows (upstream and downstream) as it is with managing the flow of products and services (see Figure 7.3).

OPERATIONS PRINCIPLE

The supply chain concept applies to the internal relationships between processes as well as the external relationships between operations.

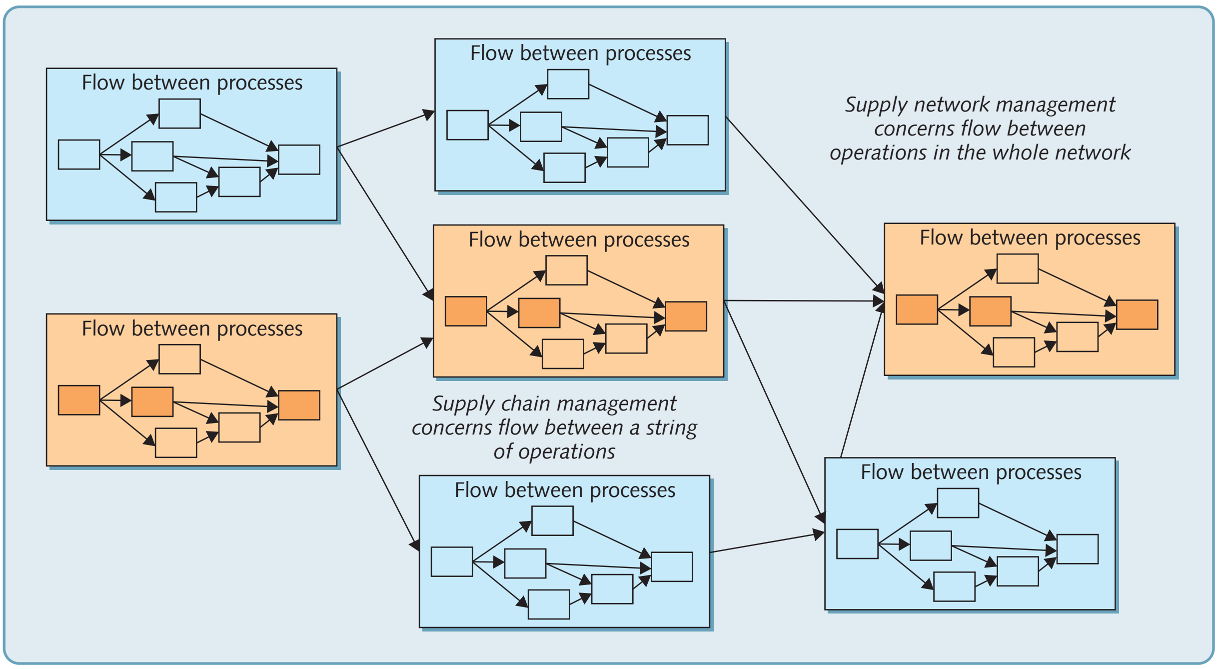

Figure 7.2 **Supply chain management is concerned with managing the flow of materials and information between a string of operations that form the strands or 'chains' of a supply network**

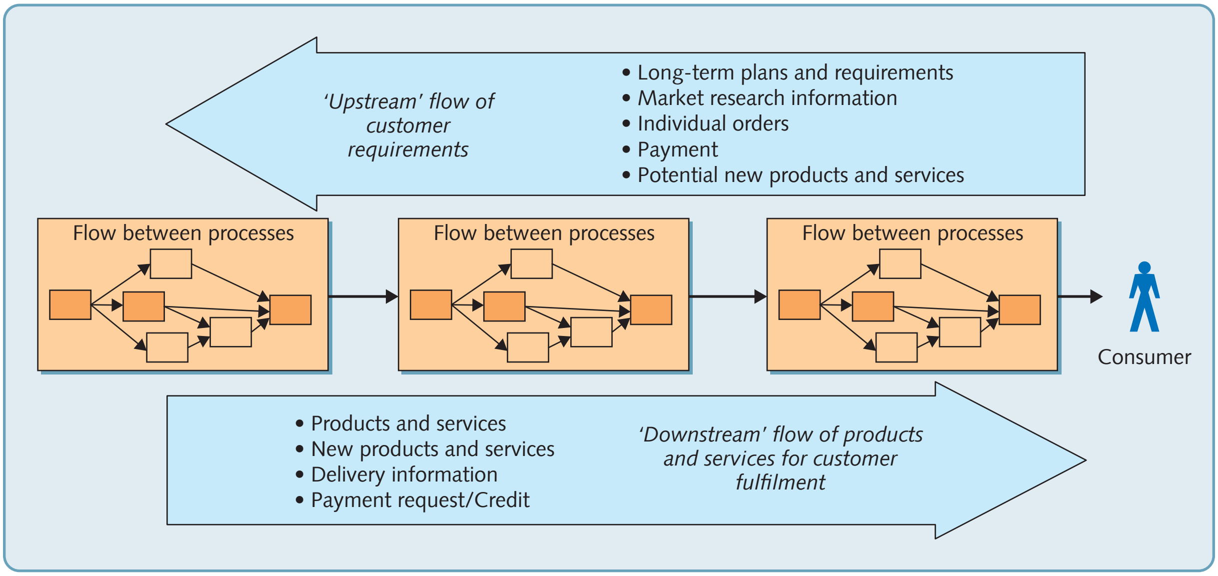

Figure 7.3 **Supply chains management is concerned with the flow of information as well as the flow of products and services**

Internal and external supply chains

Although we often describe supply chains as an interconnection of 'organisations', this does not necessarily mean that these 'organisations' are distinctly separate entities belonging to and managed by different owners. In Chapter 1 we pointed out how the idea of networks can be applied, not just at the supply network level of 'organisation-to-organisation' relationships, but also at the 'process-to-process' within-operation level and even at the 'resource-to-resource' process level. Also in Chapter 1 we introduced the idea of internal customers and suppliers. Put these two related ideas together and one can understand how some of the issues that we will be discussing in the context of 'organisation-to-organisation' supply chains can also provide insight into internal 'process-to-process' supply chains.

OPERATIONS PRINCIPLE

The supply chain concept applies to the internal relationships between processes as well as the external relationships between operations.

Tangible and intangible supply chains

Almost all the books, blogs and articles on supply chain management deal with relationships between what we called 'material transformation' operations; that is, operations concerned with the creation, movement, storage or sale of physical products. But, as we mentioned in Chapter 4, the idea of supply networks (and therefore, supply chains) applies equally to operations with largely or exclusively intangible inputs and outputs; such as financial services, retail shopping malls, insurance providers, healthcare operations, consultants, universities, and so on. All these operations have suppliers and customers, they all purchase services, they all have to choose how they get their services to consumers, in other words they all have to manage their supply chains. Of course, all supply chains, even ones that transform physical items, have service elements – again referring back to Chapter 1, most operations supply a mixture of products *and* services. So, although we, like most authors, focus primarily on supply chains that transform physical items, remember that many of the ideas discussed also have relevance for 'pure' services that supply intangibles.

OPERATIONS PRINCIPLE

The supply chain concept applies to non-physical flow between operations and processes as well as physical flows.

Why is supply chain management important?

The end consequence of what customers experience when they purchase a product or service is the result, not of an individual operation, but of a whole supply chain of operations. It is this whole supply chain that will influence final quality, speed of response, dependability, flexibility and cost of final supply. So the reputation and commercial success of an individual company depends fundamentally on its position in, and relationship with, the other operations in its supply chain. And with firms increasingly outsourcing some of its activities, especially if they have taken advantage of globalised supply chains, the coordination of supply has become particularly important. Just look what happens when incompetence or natural disasters disrupt supply chains. No operations manager can afford to ignore it.

When an earthquake off Japan's eastern coast caused a tsunami that killed thousands of people and created a meltdown at a nearby nuclear power plant, it necessitated huge evacuations and nationwide power shortages. The effect on global supply networks was immediate and drastic. Sony Corporation shut down some of its operations in Japan because of the ongoing power shortages. Toyota suspended production at most of its Japanese plants, reduced and then suspended output from its North American and European operations. Nissan said it would be suspending its UK production for three days at the end of the month due to a shortfall of parts from Japan. Honda announced that it was halving production at its factory in Swindon in the South of the UK. The lesson here is that in most modern supply chains, individual operations are often so interdependent that operations managers must broaden their focus to include the whole chain, of which they are a part.

OPERATIONS PRINCIPLE

Supply chain management is important to all operations managers because it impacts all aspects of commercial performance.

The following two examples of supply chain management illustrate some of the issues that are discussed in this chapter.

EXAMPLE

Ocado[1]

The UK grocery market is tough. Giant retailers battle for increasingly cost-conscious customers who also demand quality and service. So, when Mark Richardson, the head of Ocado's Technology operation was promoted to the newly created role of Operations Director, it was greeted as an announcement that reinforced the importance of Ocado's high-tech operations processes. It was also seen as reflecting the urgent need for Ocado to get its state-of-the-art distribution centre operating at full efficiency. The company's on-time delivery performance, although still better than its rivals, had been slipping in the months prior to the announcement. This had been an unusual experience for Ocado, which was the only dedicated online supermarket in the UK and the largest dedicated online supermarket in the world. It had succeeded in reshaping the final 'business-to-consumer' configuration of the traditional food supply chain in its UK home market, and in the process had become one of the most successful online grocers in the world.

But it was not the first. A few years before Ocado started, an internet grocer called Webvan began trading in California. It gained considerable publicity and more than a billion US dollars from backers wanting to join in what promised to be the exciting new world of online retailing. However, it proved far more difficult than Webvan's management and investors thought to make a totally new form of supply chain work. Within a couple of years it filed for bankruptcy protection. Yet Ocado has thrived.

One of its first decisions was to enter into a branding and sourcing arrangement with Waitrose, a leading high-quality UK supermarket. Just as important, it developed a supply process that provided both relative efficiency and high levels of service. (A typical Ocado delivery has a lower overall carbon footprint than walking to your local supermarket.) Most online grocers fulfil web orders by gathering goods from the shelf of a local supermarket and then loading them in a truck for delivery. By contrast, Ocado offers 'doorstep' delivery directly from its distribution centres. The orders are centrally picked from a single, state-of-the-art, highly automated warehouse (the customer fulfilment centre or 'CFC'). This is a space the size of ten football pitches, in which a 15 km system of conveyor belts handles upwards of 8,000 grocery containers an hour, which are then shipped to homes, mainly in the southern part of the UK. The largely automated picking process, which was developed by its own software engineers, allows the company to pick and prepare groceries for delivery up to seven times faster than its rivals. Also fresh or perishable items that are prepared centrally will have more 'shelf life'. Ocado's food waste, at 0.3 per cent of sales, is the lowest in the industry. The structural advantage of this supply arrangement means that 99 per cent of all orders are fulfilled accurately. Just as important as the physical distribution to the customers' door is the ease of using the company's website (Ocado.com) and the convenience of booking a delivery slot. Ocado offer reliable one-hour, next-day timeslots in an industry where two-hour timeslots prevail. This is made possible thanks, again, to the centralised model and world-class processes, systems and controls. The company say that its website is designed to be simple to use and intuitive. Smart lists personalised to each customer offer prompts and ideas so that the absence of any in-store inspiration becomes irrelevant. For a pre-registered customer, a weekly shop can be completed in less than five minutes. The site also has an extensive range of recipes, including some as video and ideas such as craft activities and lunchbox fillers. Ocado makes a conscious effort to recruit people with customer service skills and then train them as drivers, rather than vice versa. Drivers, known as Customer Service Team Members, are paid well above the industry norm and are empowered to process refunds and deal with customer concerns on the doorstep. Yet although as many as one million separate items are picked for individual customer orders every day, there are fewer than 80 mistakes.

But Ocado's supply system has its vulnerabilities. When it decided to increase the capacity of this distribution centre its construction started to run behind schedule. The automated systems had suffered from teething problems, capacity had been restricted, delivery performance adversely influenced and hiring the extra staff to handle orders had affected profitability. Because Ocado operates what it calls its 'hub and spoke' supply system; with its central CFC (hub) serving regional (spoke) distribution points it is particularly vulnerable to disruption at its 'hubs'.

EXAMPLE

Does Apple really have the best supply chain in the world?[2]

Every time Apple launches a new product it is worldwide news. They attract the type of attention and publicity that only the most famous celebrities can match. And it's similar (but lower key

maybe) in the world of operations management, where Apple is often credited with being the 'best supply chain in the world'. It is a title that is especially impressive when one considers the complexity of Apple's products and the company's strategy of frequently introducing new, and technologically advanced, products. Yet there is a connection between Apple's products and its supply chain. Innovative products that combine advanced functionality, a fast ramp-up of manufacturing capacity and customers that have a near-obsessive interest in the detail of design will typically require innovative approaches to the development of their supply chains. It integrates all its research and development, marketing, purchasing, outsourced manufacturing and logistics functions together, facilitating the detailed advanced planning that accelerated new product introduction requires, sometimes by acquiring exclusive rights from its suppliers to secure strategic raw materials and components. Early in its history the company established a formalised list of expectations for suppliers and quickly moved on to creating exclusivity agreements in exchange for volume guarantees. The relationship with suppliers is vital. Apple can use its financial muscle to guarantee sufficient supply capacity by placing large pre-orders with suppliers, which also prevents competitors from gaining access to the same manufacturing resources. Apple's supply chain is structured to give the company maximum visibility and, essentially, control over the design and nature of its products, right down to the smallest components. This focus on control was reinforced when Apple started to refine its relationships with suppliers. Conventional outsourcing usually expects the outsourcer to arrange for its own supply (Apple's suppliers). But, according to the *Wall Street Journal* (that keeps a close watch on Apple's activities) Apple takes greater control over the procurement of components for its suppliers. To do this, the company hired hundreds of engineers and supply chain managers to its staff in Shanghai and Taipei to help its suppliers.

But as well as advantages, extensive outsourcing of production does not come without risk. It exposes any company to disruption caused by anything from natural disasters to changes in international trade agreements. There are also reputational risks. Apple has suffered from criticism of the work practices in its (or rather its suppliers') extended supply chain in Asia, including a large number of suicides at a giant factory in China of its major contract manufacturer Foxconn. This explains Apple's extensive efforts to report on what is happening in terms of improvements at its suppliers. Every year they issue a report on their progress that some see as a model for other companies. '*We audit deep into our supply chain and hold our suppliers accountable to some of the industry's strictest standards,*' Apple says. '*In fact, we care as much about how our products are made as we do about how they're designed.*' Apple has also been criticised for its reliance on a relatively small number of key suppliers, what has been described as 'industrial-scale supply chain risk'.

What do they have in common?

The first lesson from these two companies is that they both take supply chain management extremely seriously. Furthermore, they both understand that, no matter how good individual operations or processes are, a business's overall performance is a function of the whole chain, of which it is a part. That is why both of these companies put so much effort into managing the whole chain. This does not mean that they adopt the same, or even similar, approaches to supply chain management. Each has a slightly different set of priorities. Ocado not so much *has* a supply chain, it *is* a supply chain in the sense that its entire purpose is supply. Moreover, it is making a business of supply using a relatively new concept (internet-based customer ordering) where other businesses have failed. It therefore needs to keep its costs under strict control if it is to avoid the fate of its predecessors like Webvan. Apple, on the other hand, needs a responsive supply chain so that its high-tech, high-design value products can have up-to-the-minute impact for its tech-savvy customers. It needs agility from its supply chain so that it can keep close to its market. It also needs to ensure (preferably exclusive) supply of state-of-the-art components, while at the same time reducing supply risk. So, although they are very different businesses with different supply chain objectives, the commonality between them is that both see the way they configure and manage their supply chains as a source of innovation. Both are doing something different from their competitors, both are innovating through their supply chain management. Both have a clear idea of what they want to be and both understand the importance of understanding customers as the starting point of successful supply chain management. Although the examples emphasised Ocado's downstream customer relationships and Apple's total supply chain (and product development) relationships, the common theme is the importance of investing in a supply perspective. In addition, both companies have invested in mechanisms for communicating along the supply chain and coordinating material and information flows. The rest of this chapter is structured around these three main issues: clarifying supply chain objectives; supply chain relationships, both with suppliers and with customers; and controlling and coordinating flow.

DIAGNOSTIC QUESTION

Are supply chain objectives clear?

All supply chain management shares one common, and central, objective – to satisfy the end customer. All stages in a chain must eventually include consideration of the final customer, no matter how far an individual operation is from the end customer. When a customer decides to make a purchase, he or she triggers action back along the whole chain. All the businesses in the supply chain pass on portions of that end customer's money to each other, each retaining a margin for the value it has added. Each operation in the chain should satisfy its own customer, but also make sure that eventually the end customer is also satisfied.

For a demonstration of how end customer perceptions of supply satisfaction can be very different from that of a single operation, examine the customer 'decision tree' in Figure 7.4. It charts the hypothetical progress of a hundred customers requiring service (or products) from a business (for example, a printer requiring paper from an industrial paper stockist). Supply performance, as seen by the core operation (the warehouse), is represented by the shaded part of the diagram. It has received 20 orders, 18 of which were 'produced' (shipped to customers) as

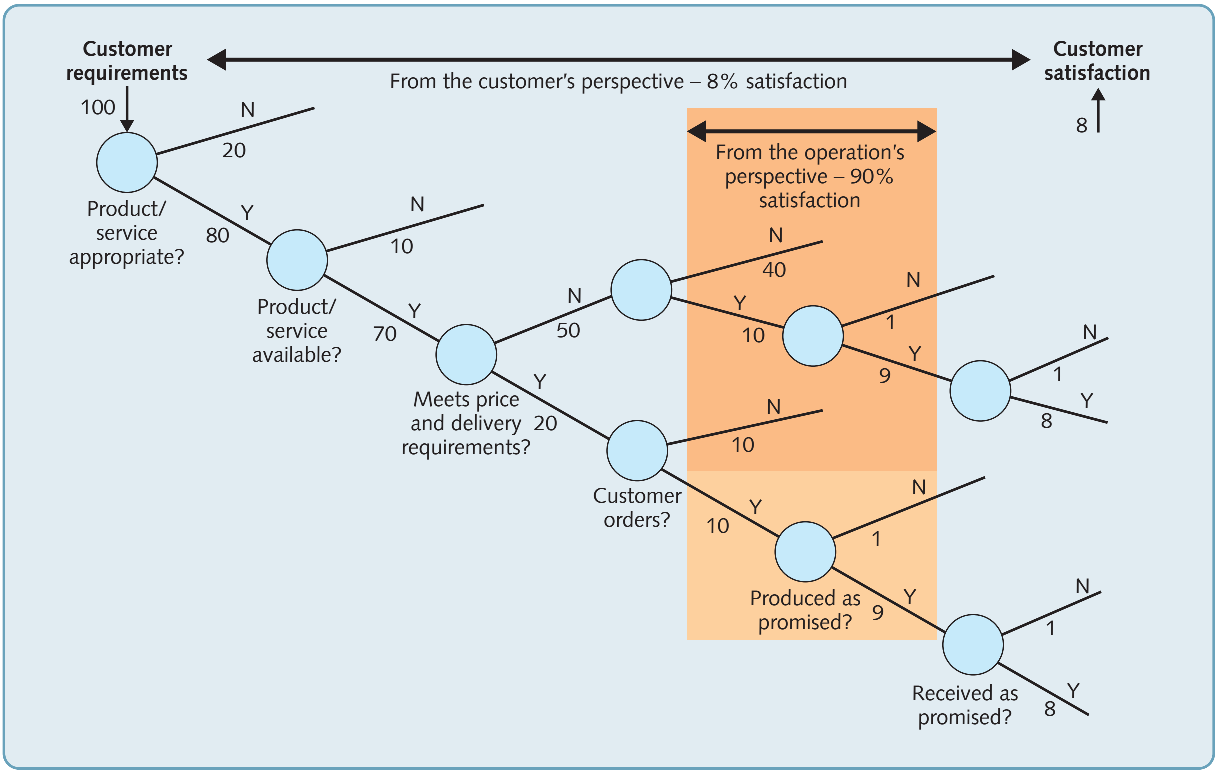

Figure 7.4 Taking a customer perspective of supply performance can lead to very different conclusions

promised (on time, and in full). However, originally 100 customers may have requested service, 20 of whom found the business did not have appropriate products (did not stock the right paper), 10 of whom could not be served because the products were not available (out of stock), 50 of whom were not satisfied with the price and/or delivery (of whom 10 placed an order notwithstanding). Of the 20 orders received, 18 were produced as promised (shipped) but two were not received as promised (delayed or damaged in transport). So what seems a 90 per cent supply performance is in fact an 8 per cent performance from the customer's perspective.

This is just one operation in a whole network. Include the cumulative effect of similar reductions in performance for all the operations in a chain, and the probability that the end customer is adequately served could become remote. The point here is not that all supply chains have unsatisfactory supply performances (although most supply chains have considerable potential for improvement). Rather it is that the performance both of the supply chain as a whole, and its constituent operations, should be judged in terms of how all end customer needs are satisfied.

OPERATIONS PRINCIPLE

The performance of an operation in a supply chain does not necessarily reflect the performance of the whole supply chain

Supply chain objectives

The objective of supply chain management is to meet the requirements of end customers by supplying appropriate products and services when they are needed at a competitive cost. Doing this requires the supply chain to achieve appropriate levels of the five operations performance objectives: quality, speed, dependability, flexibility and cost.

Quality

The quality of a product or service when it reaches the customer is a function of the quality performance of every operation in the chain that supplied it. The implication of this is that errors in each stage of the chain can multiply in their effect on end customer service (if each of seven stages in a supply chain has a 1 per cent error rate, only 93.2 per cent of products or services will be of good quality on reaching the end customer (i.e. 0.997)). This is why it is only when every stage takes some responsibility for its own *and its suppliers'* performance that a supply chain can achieve high end customer quality.

Speed

This has two meanings in a supply chain context. The first is how fast customers can be served, (the elapsed time between a customer requesting a product or service and receiving it in full), an important element in any business's ability to compete. However, fast customer response can be achieved simply by over-resourcing or over-stocking within the supply chain. For example, very large stocks in a retail operation can reduce the chances of stock-out to almost zero, so reducing customer waiting time virtually to zero. Similarly, an accounting firm may be able to respond quickly to customer demand by having a very large number of accountants on standby waiting for demand that may (or may not) occur. An alternative perspective on speed is the time taken for goods and services to move through the chain. For example, products that move quickly down a supply chain from raw material suppliers through to retailers will spend little time as inventory because to achieve fast throughput time, material cannot dwell for significant periods as inventory. This in turn reduces the working capital requirements and other inventory costs in the supply chain, so reducing the overall cost of delivering to the end customer. Achieving a balance between speed as responsiveness to customers' demands and speed as fast throughput (although they are not incompatible) will depend on how the supply chain chooses to compete.

Dependability

Dependability in a supply chain context is similar to speed in so much as one can almost guarantee 'on-time' delivery by keeping excessive resources, such as inventory, within the chain. However, dependability of throughput time is a much more desirable aim because it reduces uncertainty within the chain. If the individual operations in a chain do not deliver as promised on time, there will be a tendency for customers to over-order, or order early, in order to provide some kind of insurance against late delivery. The same argument applies if there is uncertainty regarding the *quantity* of products or services delivered. This is why delivery dependability is often measured as 'on time, in full' in supply chains.

Flexibility

In a supply chain context, flexibility is usually taken to mean the chain's ability to cope with changes and disturbances. Very often this is referred to as supply chain agility. The concept of agility includes previously discussed issues, such as focusing on the end customer and ensuring fast throughput and responsiveness to customer needs. But, in addition, agile supply chains are sufficiently flexible to cope with changes, either in the nature of customer demand, or in the supply capabilities of operations within the chain.

Cost

In addition to the costs incurred within each operation to transform its inputs into outputs, the supply chain as a whole incurs additional costs that derive from each operation in a chain doing business with each other. These transaction costs may include such things as the costs of finding appropriate suppliers, setting up contractual agreements, monitoring supply performance, transporting products between operations, holding inventories, and so on. Many of the recent developments in supply chain management, such as partnership agreements or reducing the number of suppliers, are an attempt to minimise transaction costs.

Sustainability

Any organisation that subscribes to environmentally responsible practices itself will want to make sure that it purchases its input products and services from suppliers that are similarly responsible. The concept is called 'sustainable procurement'. It is a process whereby organisations 'meet their needs for goods, services, works and utilities in a way that achieves value for money on a whole life basis in terms of generating benefits not only to the organisation, but also to society and the economy, while minimising damage to the environment'.[3] Operations that try to follow sustainable purchasing often recommend practices such as: buying from local vendors where possible; sourcing supplies from suppliers with ethical practices; choosing environmentally friendly products and services; using minimal packaging; transporting products by ground transport rather than air, and so on.

EXAMPLE

North Face's sustainable purchasing[4]

Few outdoor clothing brands have had the impact of The North Face® since it was founded over 40 years ago in San Francisco. It is part of the VF Corporation, a $9 billion giant that dominates the leisurewear garment market. According to the influential WebTogs blog, '*one of the brilliant things about The North Face jackets is that no matter whether you're off to battle with some sub-zero temperatures or just to chill out and take a load off there is always a North Face jacket to accompany you and keep you warm and snug*'. Its range of high-performance outdoor apparel, equipment and footwear has developed a reputation for durability, fashionable styling, and, increasingly, sustainable sourcing of its materials. VF Corporation's claim is that they 'responsibly manage the industry's most efficient and complex supply chain, which spans multiple geographies, product categories and distribution channels'. In particular, The North Face's commitment to sustainability, they say, comes from a desire to protect the natural places associated with how and where their products are used and from their concern about the effects of climate change. As a sign of their determination to pursue sustainable purchasing, they have partnered with the independent bluesign®

standard, a Swiss-based organisation whose objective is to tackle sustainability at its roots and exclude substances and practices that are potentially hazardous to human health or the environment from all processes in the garment supply chain. To be considered as a supplier to The North Face any factory must meet the rigorous, independent bluesign® standards.

Should supply chains be lean or agile?

A distinction is often drawn between supply chains that are managed to emphasise supply chain efficiency (lean supply chains), and those that emphasise supply chain responsiveness and flexibility (agile supply chains). These two modes of managing supply chains are reflected in an idea proposed by Professor Marshall Fisher of Wharton Business School;[5] that supply chains serving different markets should be managed in different ways. Even companies that have seemingly similar products or services, in fact, may compete in different ways with different products. For example, shoe manufacturers may produce classics that change little over the years, as well as fashion shoes that last only one season. Chocolate manufacturers have stable lines that have been sold for 50 years, but also product 'specials' associated with an event or film release, the latter selling only for a matter of months. Hospitals have routine 'standardised' surgical procedures such as cataract removal, but also have to provide emergency post-trauma surgery. Demand for the former products will be relatively stable and predictable, but demand for the latter will be far more uncertain. Also, the profit margin commanded by the innovative product will probably be higher than that of the more functional product. However, the price (and therefore the margin) of the innovative product may drop rapidly once it has become unfashionable in the market.

The supply chain policies that are seen to be appropriate for functional products and innovative products are termed efficient (or lean) and responsive (or agile) supply chain policies, respectively. Efficient supply chain policies include keeping inventories low, especially in the downstream parts of the network, so as to maintain fast throughput and reduce the amount of working capital tied up in the inventory. What inventory there is in the network is concentrated mainly in the manufacturing operation, where it can keep utilisation high and therefore manufacturing costs low. Information must flow quickly up and down the chain from retail outlets back up to the manufacturer so that schedules can be given the maximum amount of time to adjust efficiently. The chain is then managed to make sure that products flow as quickly as possible down the chain to replenish what few stocks are kept downstream.

OPERATIONS PRINCIPLE
Supply chains with different end objectives need managing differently.

By contrast, responsive supply chain policy stresses high service levels and responsive supply to the end customer. The inventory in the network will be deployed as closely as possible to the customer. In this way, the chain can still supply even when dramatic changes occur in customer demand. Fast throughput from the upstream parts of the chain will still be needed to replenish downstream stocks. But those downstream stocks are needed to ensure high levels of availability to end customers. Figure 7.5 illustrates how the different supply chain policies match the different market requirements implied by functional and innovative products.

OPERATIONS PRINCIPLE
'Functional' products require lean supply chain management; 'innovative' products require agile supply chain management.

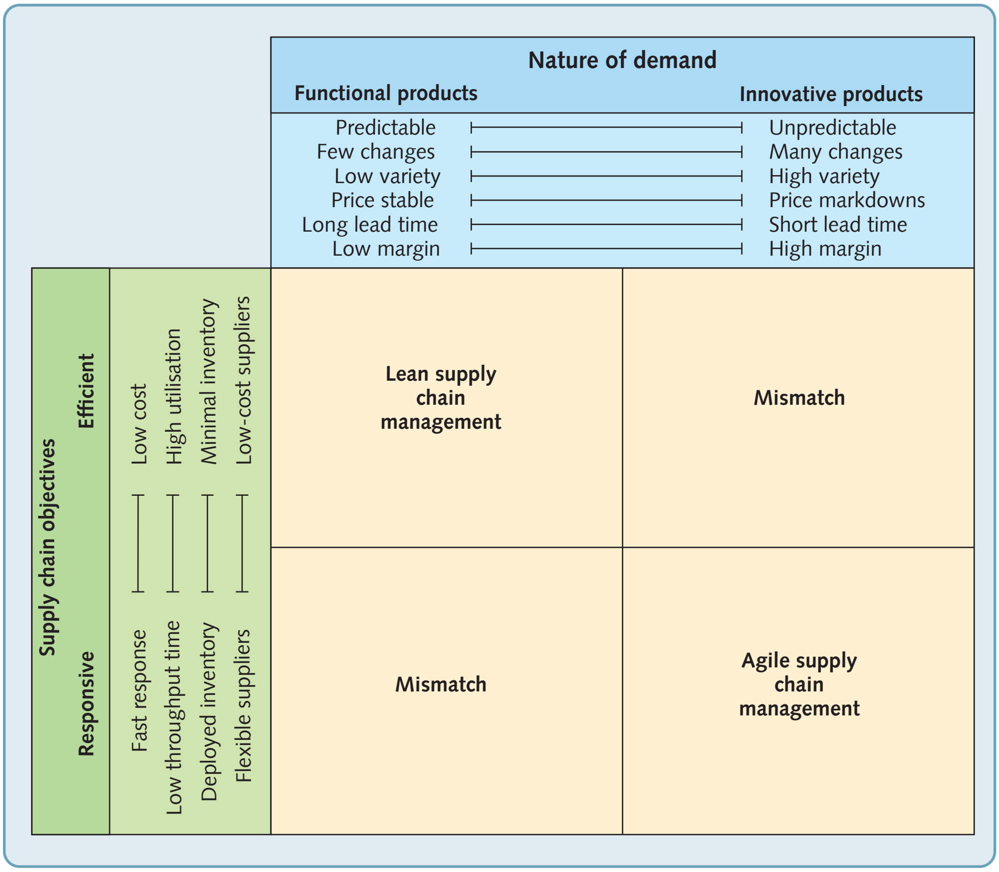

Figure 7.5 Matching the operations resources in the supply chain with market requirements

Source: Adapted from Fisher, M.C. (1997) 'What is the right supply chain for your product?' *Harvard Business Review*, March–April, 105–116.

DIAGNOSTIC QUESTION

How should supply chain relationships be managed?

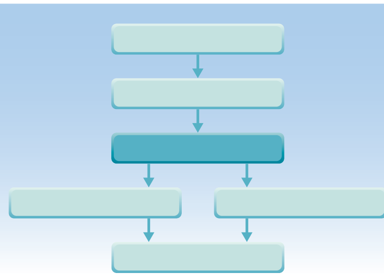

The 'relationship' between operations in a supply chain is the basis on which the exchange of products, services, information and money is conducted. Managing supply chains is about managing relationships, because relationships influence the smooth flow between operations and processes. Different forms of relationship will be appropriate in different circumstances. An obvious but important factor in determining the importance of relationships to any operation is the extent to which they outsource their activities. In Chapter 4 we distinguished between non-vertically integrated operations that outsource almost all activities, and vertically integrated operations that outsource almost nothing. Only extremely vertically integrated businesses are able to ignore the question of how to manage customer–supplier relationships (because they do everything themselves). Initially, we can examine this question by describing two contrasting 'pure' arrangements: pure *contractual*, market-based, transactional relationships; and close, longer term, pure *partnership* relationships. However, it is better to think of these as the two basic ingredients of any supply arrangement. Whatever arrangement with its suppliers a firm chooses to take, it can be described by the balance between contracts and partnerships.

Contract-based 'transactional' relationships

Contract-based, transactional relationships involve purchasing goods and services in a 'pure' market fashion, often seeking the 'best' supplier every time it is necessary to make a purchase. Each transaction effectively becomes a separate decision. The relationship may be short term, with no guarantee of further trading between the parties once the goods or services are delivered and payment is made.[6] The *advantages* of contract-based 'transactional' relationships are usually seen as follows:

- They maintain competition between alternative suppliers. This promotes a constant drive between suppliers to provide best value.
- A supplier specialising in a small number of products or services, but supplying them to many customers, can gain natural economies of scale, enabling the supplier to offer the products and services at a lower price than if customers performed the activities themselves on a smaller scale.
- There is inherent flexibility in outsourced supplies. If demand changes, customers can simply change the number and type of suppliers; a faster and cheaper alternative to redirecting internal activities.
- Innovations can be exploited no matter where they originate. Specialist suppliers are more likely to come up with innovations that can be acquired faster and cheaper than developing them in-house.

There are, however, *disadvantages* in buying in a totally contractual manner:

- Suppliers owe little loyalty to customers. If supply is difficult, there is no guarantee of receiving supply.
- Choosing who to buy from takes time and effort. Gathering sufficient information and making decisions continually are, in themselves, activities that need to be resourced.

Short-term contractual relationships of this type may be appropriate when new companies are being considered as more regular suppliers, or when purchases are one-off or very irregular (for example, the replacement of all the windows in a company's office block would typically involve this type of competitive-tendering market relationship).

Long-term 'partnership' relationships

Partnership relationships in supply chains are sometimes seen as a compromise between vertical integration on the one hand (owning the resources which supply you) and transactional relationships on the other. Partnership relationships are defined as:[7] '. . . *relatively enduring inter-firm cooperative agreements, involving flows and linkages that use resources and/or governance structures from autonomous organizations, for the joint accomplishment of individual goals linked to the corporate mission of each sponsoring firm.*' This means that suppliers and customers are expected to cooperate, even to the extent of sharing skills and resources, to achieve joint benefits beyond those they could have achieved by acting alone. At the heart of the concept of partnership lies the issue of the *closeness* of the relationship. Partnerships are close relationships, the degree of which is influenced by a number of factors, as follows:

- *Sharing success* – both partners jointly benefit from the cooperation rather than manoeuvring to maximise their own individual contribution.
- *Long-term expectations* – relatively long-term commitments, but not necessarily permanent ones.

- *Multiple points of contact* – communication not restricted to formal channels, but may take place between many individuals in both organisations.
- *Joint learning* – a relationship commitment to learn from each other's experience.
- *Few relationships* – a commitment on the part of both parties to limit the number of customers or suppliers with whom they do business.
- *Joint coordination of activities* – fewer relationships allow joint coordination of activities such as the flow of materials or service, payment, and so on.
- *Information transparency* – confidence is built through information exchange between the partners.
- *Joint problem solving* – jointly approaching problems can increase closeness over time.
- *Trust* – probably the key element in partnership relationships. In this context, trust means the willingness of one party to relate to the other on the understanding that the relationship will be beneficial to both, even though that cannot be guaranteed. Trust is widely held to be both the key issue in successful partnerships, but also, by far, the most difficult element to develop and maintain.

Which type of relationship?

OPERATIONS PRINCIPLE

All supply chain relationships can be described by the balance between their 'contractual' and 'partnership' elements.

It is very unlikely that any business will find it sensible to engage exclusively in one type of relationship or another. Most businesses will have a portfolio of, possibly, widely differing relationships. Also, there are degrees to which any particular relationship can be managed on a transactional or partnership basis. The real question is: Where, on the spectrum from transactional to partnership, should each relationship be positioned? And, while there is no simple formula for choosing the 'ideal' form of relationship in each case, there are some important factors that can sway the decision. The most obvious issue will concern how a business intends to compete in its market place. If price is the main competitive factor then the relationship could be determined by which approach offers the highest potential savings. On the one hand, market-based contractual relationships could minimise the actual price paid for purchased products and services, while partnerships could minimise the transaction costs of doing business. If a business is competing primarily on product or service innovation, the type of relationship may depend on where innovation is likely to happen. If innovation depends on close collaboration between supplier and customer, partnership relationships are needed. On the other hand, if suppliers are busily competing to out-do each other in terms of their innovations, and especially if the market is turbulent and fast growing (as with many software and internet-based industries), then it may be preferable to retain the freedom to change suppliers quickly using market mechanisms. However, if markets are very turbulent, partnership relationships may reduce the risks of being unable to secure supply.

The main differences between the two ends of this relationship spectrum concerns whether a customer sees advantage in long-term or short-term relationships. Contractual relationships can be either long or short term, but there is no guarantee of anything beyond the immediate contract. They are appropriate when short-term benefits are important. Many relationships and many businesses are best by concentrating on the short term (especially if, without short-term success, there is no long term). Partnership relationships are by definition long term. There is a commitment to work together over time to gain mutual advantage. The concept of mutuality is important here. A supplier does not become a 'partner' merely by being called one. True partnership implies mutual benefit, and often mutual sacrifice. Partnership means giving up some freedom of action in order to gain something more beneficial over the long term. If it is not in the culture of a business to give up some freedom of action, it is very unlikely to ever make a success of partnerships. Opportunities to develop relationships can be limited by the structure of the market itself. If the number of potential suppliers is small, there may be few

OPERATIONS PRINCIPLE

True 'partnership' relationships involve mutual sacrifice as well as mutual benefit.

opportunities to use market mechanisms to gain any kind of supply advantage and it would probably be sensible to develop a close relationship with at least one supplier. On the other hand, if there are many potential suppliers, and especially if it is easy to judge the capabilities of the suppliers, contractual relationship are likely to be best.

DIAGNOSTIC QUESTION

How is the supply side managed?

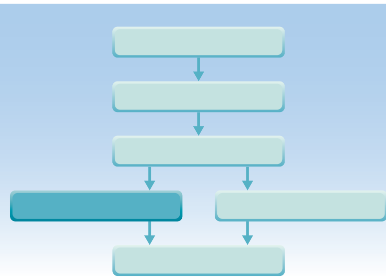

The ability of any process or operation to produce outputs is dependent on the inputs it receives. So good supply management is a necessary (but not sufficient) condition for effective operations management in general. It involves three main activities: selecting appropriate suppliers; planning and controlling the ongoing supply activity; and developing and improving suppliers' capabilities. All three activities are usually the responsibility of the purchasing or procurement function within the business. Purchasing should provide a vital link between the operation itself and its suppliers. They should understand the requirements of all the processes within their own operation and also the capabilities of the suppliers who could potentially provide products and services for the operation.

Supplier selection

Choosing appropriate suppliers should involve trading off alternative attributes. Rarely are potential suppliers so clearly superior to their competitors that the decision is self-evident. Most businesses find it best to adopt some kind of supplier 'scoring' or assessment procedure. This should be capable of rating alternative suppliers in terms of factors such as the following:

- range of products or services provided
- quality of products or services
- responsiveness
- dependability of supply
- delivery and volume flexibility
- total cost of being supplied
- ability to supply in the required quantity.

In addition, there are likely to be less quantifiable or longer term factors that will need to be taken into consideration. These may include the following:

- potential for innovation
- ease of doing business
- willingness to share risk
- long-term commitment to supply
- ability to transfer knowledge as well as products and services.

Choosing suppliers should involve evaluating the relative importance of all these factors. So, for example, a business might choose a supplier who, although more expensive than alternative suppliers, has an excellent reputation for on-time delivery, because that is more appropriate to the way the business competes itself, or because the high level of supply dependability allows the business to hold lower stock levels, which may even save costs overall. Other trade-offs may be more difficult to calculate. For example, a potential supplier may have high levels of technical capability, but may be financially weak, with a small but finite risk of going out of business. Other suppliers may have little track record of supplying the products or services required, but show the managerial talent and energy for potential customers to view developing a supply relationship as an investment in future capability. But in order to make sensible trade-offs it is important to assess four basic capabilities:

OPERATIONS PRINCIPLE

Supplier selection should reflect overall supply chain objectives.

- *Technical capability* – the product or service knowledge to supply to high levels of specification.
- *Operations capability* – the process knowledge to ensure consistent, responsive, dependable and reasonable cost supply.
- *Financial capability* – the financial strength to fund the business in the short and long terms.
- *Managerial capability* – the management talent and energy to develop supply potential in the future.

Single- or multi-sourcing

A closely linked decision is whether to source each individual product or service from one, or more than one, supplier (single-sourcing or multi-sourcing). Some of the advantages and disadvantages of single- and multi-sourcing are shown in Table 7.1.

It may seem as though companies who multi-source do so exclusively for their own short-term benefit. However, this is not always the case: multi-sourcing can have an altruistic motive, or at least one that brings benefits to both supplier and purchaser in the long term. For example, Robert Bosch GmbH, the German automotive components manufacturer and distributor, at one time required that subcontractors do no more than 20 per cent of their total business with them.[8] This was to prevent suppliers becoming too dependent on them. The purchasing organisation could then change volumes up and down without pushing the supplier into bankruptcy. However, despite these perceived advantages, there has been a trend for purchasing functions to reduce their supplier base in terms of numbers of companies supplying any one part or service, mainly because it reduces the costs of transacting business.

Table 7.1 Advantages and disadvantages of single- and multi-sourcing

	Single-sourcing	*Multi-sourcing*
Advantages	Potentially better quality because of more supplier quality assurance possibilities Strong relationships that are more durable Greater dependency encourages more commitment and effort Better communication Easier to cooperate on new product/service development More economies of scale Higher confidentiality	Purchaser can drive price down by competitive tendering Can switch sources in case of supply failure Wide sources of knowledge and expertise to tap
Disadvantages	More vulnerable to disruption if a failure to supply occurs Individual supplier more affected by volume fluctuations Supplier might exert upward pressure on prices if no alternative supplier is available	Difficult to encourage commitment supplier Less easy to develop effective supplier quality assurance More effort needed to communicate Suppliers less likely to invest in new processes More difficult to obtain economies of scale

Purchasing, the internet and e-commerce

For some years, electronic means have been used by businesses to confirm purchased orders and ensure payment to suppliers. The rapid development of the internet, however, opened up the potential for far more fundamental changes in purchasing behaviour. Partly this was as the result of supplier information made available through the internet. Previously, a purchaser of industrial components may have been predisposed to return to suppliers who had been used before. There was inertia in the purchasing process because of the costs of seeking out new suppliers. By making it easier to search for alternative suppliers, the internet changes the economics of the search process and offers the potential for wider searches. It also changed the economies of scale in purchasing. Purchasers requiring relatively low volumes find it easier to group together in order to create orders of sufficient size to warrant lower prices. In fact, the influence of the internet on purchasing behaviour is not confined to *e-commerce*. Usually, e-commerce is taken to mean the trade that actually takes place over the internet. This is usually assumed to be a buyer visiting the seller's website, placing an order for parts and making a payment (also through the site). But the internet is also an important source of purchasing information. For every 1 per cent of business transacted directly via the internet, there may be 5 or 6 per cent of business that, at some point, involved it, probably with potential buyers using it to compare prices or obtain technical information.

One increasingly common use of internet technology in purchasing (or e-procurement as it is sometimes known) is for large companies, or groups of companies, to link their e-commerce systems into a common 'exchange'. In their more sophisticated form, such an exchange may be linked into the purchasing companies' own information systems (see the explanation of ERP in Chapter 10). Many of the large automotive, engineering and petrochemical companies, for example, have adopted such an approach. An early example of this was Dow Corning's 'Xiameter' service. Dow Corning was the global market leader in silicon, a material which has a wide range of industrial applications, from clothing and computers, to cosmetics construction. Traditionally, its customers had paid top prices for pioneering technology and premium quality products, delivered with an emphasis on 'solutions-based service. However, in the early 2000s some of its larger and more sophisticated customers wanted a different kind of supply arrangement. As one customer put it; *'I don't need these services. I know I can go and buy a tanker of this fluid at a lower price. I'll buy this but I just need low price and guaranteed delivery.'* This part of their market consisted of experienced purchasers of commonly used silicone materials who wanted the lowest price and an easy way of doing business with their supplier. Their solution was to offer a 'no-frills', limited-availability service with low prices that could only be accessed on the web. This service would offer only regular products without any technical advice. They branded this service 'Xiameter' (rhymes with 'diameter'). It was limited to about 350 common silicone compounds (out of more than 7,500) ordered in high volumes. It would have a 'lean' management structure and would secure its supply from Dow Corning's manufacturing sites around the world. Customers could only place orders online (a novel approach at the time). Minimum order quantities were strictly applied. Delivery lead times were fixed by production scheduling.

Managing ongoing supply

Managing supply relationships is not just a matter of choosing the right suppliers and then leaving them to get on with day-to-day supply. It is also about ensuring that suppliers are given the right information and encouragement to maintain smooth supply and that internal inconsistency does not negatively affect their ability to supply. A basic requirement is that some mechanism should be set up that ensures the two-way flow of information between customer and supplier. It is easy for both suppliers and customers simply to forget to inform each other of internal developments that could affect supply. Customers may see suppliers as having the responsibility for ensuring appropriate supply 'under any circumstances'. Or, suppliers themselves may be reluctant to inform customers of any potential problems with supply because they see it as risking the relationship. Yet, especially if customer and supplier see themselves as 'partners', the free flow of information, and a mutually supportive tolerance of occasional problems, is the best way to ensure smooth supply. Often day-to-day supplier relationships are damaged because of internal inconsistencies. For example, one part of a business may be asking a supplier for some special service beyond the strict nature of their agreement, while another part of the business is not paying suppliers on time.[9]

Service-level agreements

Some organisations bring a degree of formality into supplier relationships by encouraging (or requiring) all suppliers to agree service-level agreements (SLAs). SLAs are formal definitions of the dimensions of service and the relationship between suppliers and the organisation. The type of issues covered by such an agreement could include response times, the range of services, dependability of service supply, and so on. Boundaries of responsibility and appropriate performance measures could also be agreed. For example, an SLA between an information systems support unit and a research unit in the laboratories of a large pharmaceutical company could define such performance measures as:

- the types of information network services which may be provided as 'standard'
- the range of special information services which may be available at different periods of the day
- the minimum 'up time', i.e. the proportion of time the system will be available at different periods of the day
- the maximum response time and average response time to get the system fully operational should it fail
- the maximum response time to provide 'special' services, and so on.

Although SLAs are described here as mechanisms for governing the ongoing relationship between suppliers and customers, they often prove to be inadequate because they are seen as being useful in setting up the terms of the relationship, but then are only used to resolve disputes. For SLAs to work effectively, they must be treated as working documents that establish the details of ongoing relationships *in the light of experience.* Used properly, they are a repository of the knowledge that both sides have gathered through working together. Any SLA that stays unchanged over time is, at the very least, failing to encourage improvement in supply.

How can suppliers be developed?

In any relationship other than pure market-based transactional relationships, it is in a customer's long-term interests to take some responsibility for developing supplier capabilities. Helping a supplier to improve not only enhances the service (and hopefully price) from the supplier, it may also lead to greater supplier loyalty and long-term commitment. This is why some particularly successful businesses (including Japanese automotive manufacturers) invest in supplier development teams, whose responsibility is to help suppliers to improve their own operations processes. Of course, committing the resources to help suppliers is only worthwhile if it improves the effectiveness of the supply chain as a whole. Nevertheless, the potential for such enlightened self-interest can be significant.

How customers and suppliers see each other[10]

One of the major barriers to supplier development is the mismatch between how customers and suppliers perceive both what is required and how the relationship is performing. Exploring potential mismatches is often a revealing exercise, both for customers and suppliers. Figure 7.6 illustrates

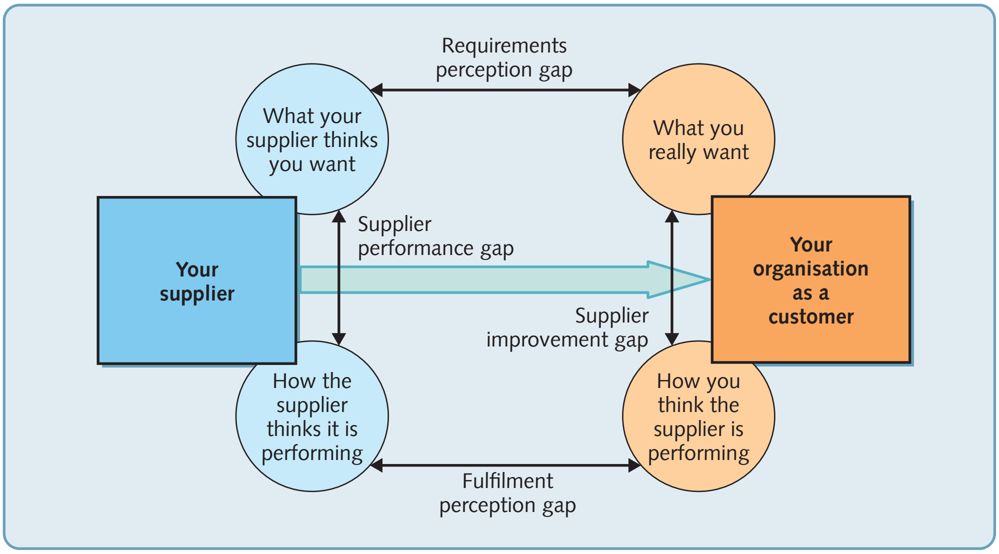

Figure 7.6 Explore the potential perception mismatches to understand supplier development needs

this. It shows that gaps may exist between four sets of ideas. As a customer you (presumably) have an idea about what you really want from a supplier. This may, or may not, be formalised in the form of a service-level agreement. But no SLA can capture everything about what is required. There may be a gap between how you as a customer interpret what is required and how the supplier interprets it. This is the *requirements perception gap*. Similarly, as a customer, you (again presumably) have a view on how your supplier is performing in terms of fulfilling your requirements. That may not coincide with how your supplier believes it is performing. This is the *fulfilment perception gap*. Both these gaps are a function of the effectiveness of the communication between supplier and customer. But there are also two other gaps. The gap between what you want from your supplier and how they are performing indicates the type of development that, as a customer, you should be giving to your supplier. Similarly, the gap between your supplier's perceptions of your needs and its performance indicates how they should initially see themselves improving their own performance. Ultimately, of course, their responsibility for improvement should coincide with their customer's views of requirements and performance.

> OPERATIONS PRINCIPLE
> ***Unsatisfactory supplier relationships can be caused by requirements and fulfilment perception gaps.***

DIAGNOSTIC QUESTION

How is the demand side managed?

The management of demand side relationships will depend partly on the nature of demand, in particular how uncertain it is. Knowing the exact demands that customers are going to require allows a supplier to plan its own internal processes in a systematic manner. This type of demand is called 'dependent' demand; it is relatively predictable because it is dependent upon some factor that is itself predictable. For example, supplying tyres to an automobile factory involves examining the manufacturing schedules in the car plant and deriving the demand for tyres from these. If

200 cars are to be manufactured on a particular day, then it is simple to calculate that 1,000 tyres will be demanded by the car plant (each car has five tyres). Because of this, the tyres can be ordered from the tyre manufacturer to a delivery schedule that is closely in line with the demand for tyres from the plant. In fact, the demand for every part of the car plant will be derived from the assembly schedule for the finished cars. Manufacturing instructions and purchasing requests will all be dependent upon this figure. Managing internal process networks when external demand is dependent is largely a matter of calculating, in as precise a way as possible, the internal consequences of demand. MRP, discussed in Chapter 10, is the best-known dependent demand approach.

But not all operations have such predictable demand. Some operations are subject to independent demand. There is a random element in demand that is virtually independent of any obvious factors. They are required to supply demand without having any firm forward visibility of customer orders. A drive-in tyre replacement service will need to manage a stock of tyres. In that sense, it is exactly the same task that faced the supplier of tyres to the car plant, but demand is very different. It cannot predict either the volume, or the specific needs of customers. It must make decisions on how many and what type of tyres to stock, based on demand forecasts and in the light of the risks it is prepared to run of being out of stock. Managing internal process networks when external demand is independent involves making 'best guesses' concerning future demand, attempting to put the resources in place to satisfy this demand, and attempting to respond quickly if actual demand does not match the forecast. Inventory planning and control, discussed in Chapter 9, is a typical approach.

Logistics services

Logistics means moving products to customers. Sometimes the term 'physical distribution management' or simply 'distribution' is used as being analogous to logistics. An important decision is how much of the logistical process of organising the movement of goods to trust to outside service providers. The extent and integration of this type of service provision is often referred to as First, Second, Third or Fourth Party Logistics (or 1PL, 2PL, 3PL, 4PL, for short). However, the distinction between the PL classifications can sometimes be blurred, with different firms using slightly different definitions:

- *First Party Logistics (1PL)* – is when, rather than outsourcing the activity, the owner of whatever is being transported organises and performs product movements themselves. For example, a manufacturing firm will deliver directly, or a retailer such as a supermarket will collect products from a supplier. The logistics activity is an entirely internal process.
- Second Party Logistics (2PL) – is when a firm decides to outsource or subcontract logistics services over a specific segment of a supply chain. It could involve a road, rail, air, or maritime shipping company being hired to transport, and if necessary store, products from a specific collection point to a specific destination.
- *Third Party Logistics (3PL)* – is when a firm contracts a logistics company to work with other transport companies to manage their logistics operations. It is a broader concept than 2PL and can involve transportation, warehousing, inventory management and even packaging or re-packaging products. Generally, 3PL involves services that are scaled and customised to a customer's specific needs.
- *Fourth Party Logistics (4PL)* – is a yet broader idea than 3PL. Accenture, the consulting group, originally used the term '4PL'. Their definition of 4PL is: 'A 4PL is an integrator that assembles the resources, capabilities, and technology of its own organisation and other organisations to design, build and run comprehensive supply chain solutions.' 4PL service suppliers pool transport capabilities, processes, technology support and coordination activities to provide customised supply chain services for part or all of a client's supply chain. 4PL firms can manage all aspects of a client's supply chain. They may act as a single interface between the client

and multiple logistics service providers, and are often separate organisational entities, founded on a long-term basis or as a joint venture between a client and one or more partners.

- *5PL?* – you guessed it, almost inevitably, some firms are selling themselves as Fifth Party Logistics providers, mainly by defining themselves as broadening the scope further to e-business.

Logistics management and the Internet of Things (IoT)

Internet-based communication has had a significant impact on physical distribution management, in particular through the adoption of the 'Internet of Things' (IoT). The IoT is a description first used at the Auto-ID Center at the Massachusetts Institute of Technology (MIT), which created a set of standards for RFID. At its simplest, an IoT is a network of physical objects (such as products, equipment, materials handling devices, trucks, etc.) that have electronics, software and sensors implanted in them, which can gather and exchange data, usually using some form of network connectivity (broadband internet, Wi-Fi, Bluetooth, etc.). Combine this with global positioning systems (GPS) that permit instantaneous tracking of trucks, materials and people, and logistics companies, warehouses, suppliers and customers can share knowledge of where goods are in the chain and where they are going next. This allows the operations within the chain to coordinate their activities more readily. It also gives the potential for some significant cost savings. For example, an important issue for transportation companies is back-loading. When the company is contracted to transport goods from A to B, its vehicles may have to return from B to A empty. Back-loading means finding a potential customer who wants their goods transported from B to A in the right timeframe. With the increase in information availability through the internet, the possibility of finding a back-load increases. Companies that can fill their vehicles on both the outward and return journeys will have significantly lower costs per distance travelled than those whose vehicles are empty for half the total journey. Similarly, internet-based technology that allows customers visibility of the progress of distribution can be used to enhance the perception of customer service. 'Track-and-trace' technologies, for example, allow package distribution companies to inform and reassure customers that their service is being delivered as promised.

Customer development

Earlier in the chapter, Figure 7.6 illustrated some of the gaps in perception and performance that can occur between customers and suppliers. The purpose then was to demonstrate the nature of supplier development. The same approach can be used to analyse the nature of requirements and performance with customers. In this case, the imperative is to understand customer perceptions, both of their requirements and their view of your performance, and feed these into your own performance improvement plans. What is less common, but can be equally valuable, is to use these gaps (shown in Figure 7.7) to examine the question of whether customer requirements and perceptions of performance are either accurate or reasonable. For example, customers may place demands on suppliers without fully considering their consequences. It may be that slight modifications in what is demanded would not inconvenience customers and yet would provide significant benefits to suppliers that could then be passed on to customers. Similarly, customers may be incompetent at measuring supplier performance, in which case the benefits of excellent supplier service will not be recognised. So, just as customers have a responsibility to help develop their own supplier's performance, in their own as well as their supplier's interests, suppliers have a responsibility to develop their customer's understanding of how supply should be managed.

OPERATIONS PRINCIPLE

Unsatisfactory customer relationships can be caused by requirement and fulfilment perception gaps.

EXAMPLE

The 80,000 kilometre journey of Wimbledon's tennis balls

The Wimbledon 'Grand Slam' tennis tournament is a quintessentially British occasion, and Slazenger, the UK sports equipment manufacturer, has been the official ball supplier for Wimbledon since 1902. Yet those balls used at Wimbledon, and the materials from which they are made,

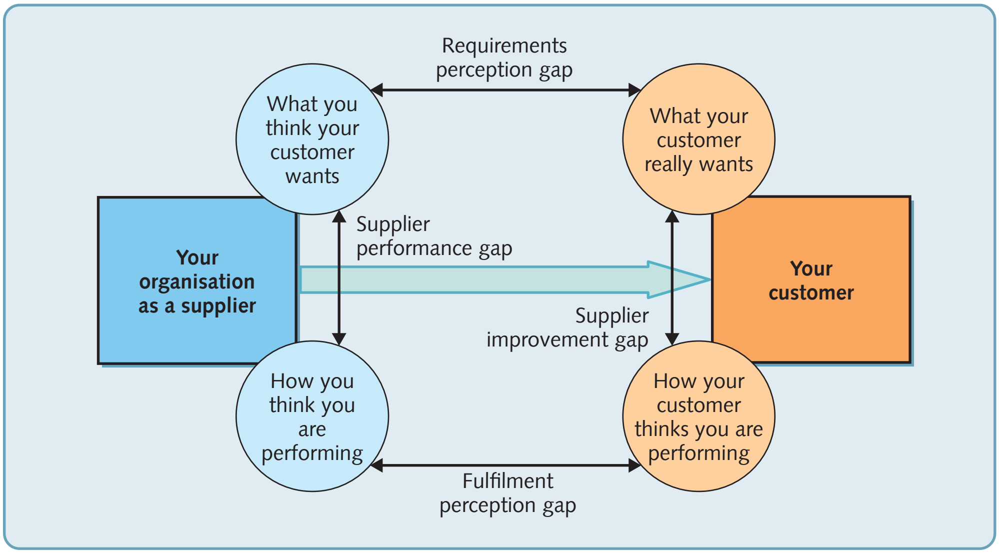

Figure 7.7 **Explore the potential perception mismatches to understand customer development needs**

will have travelled 81,385 kilometres between 11 countries and across four continents before they reach Centre Court. Dr. Mark Johnson, of Warwick Business School said: 'It is one of the longest journeys I have seen for a product.' On the face of it, travelling more than 80,000 kilometres to make a tennis ball does seem fairly ludicrous, but it just shows the global nature of production, and in the end, this will be the most cost-effective way of making tennis balls. Slazenger are locating production near the primary source of their materials in Bataan in the Philippines, where labour is also relatively low cost.

The complex supply chain is illustrated in Figure 7.8. It sees clay shipped from South Carolina in the USA, silica from Greece, magnesium carbonate from Japan, zinc oxide from Thailand, sulphur from South Korea and rubber from Malaysia to Bataan where the rubber is vulcanised – a chemical

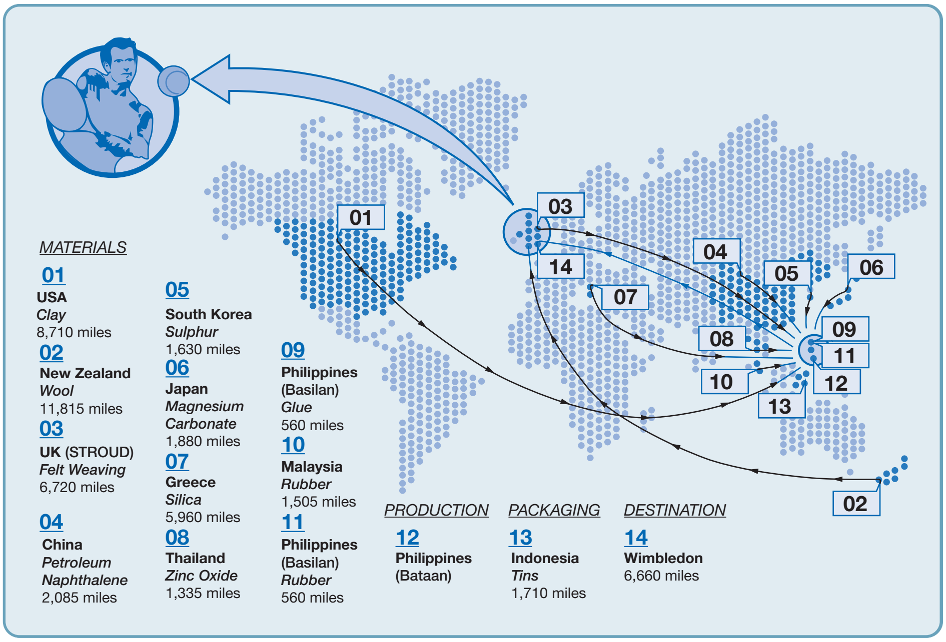

Figure 7.8 **Wimbledon's tennis balls travel over 80,000 kilometres in their supply network**

process for making the rubber more durable. Wool is then shipped from New Zealand to Stroud in the UK, where it is woven into felt and then flown back to Bataan in the Philippines. Meanwhile, petroleum naphthalene from Zibo in China and glue from Basilan in the Philippines are brought to Bataan, where Slazenger manufacture the ball. Finally, the tins, which contain the balls, are shipped in from Indonesia and once the balls have been packaged they are sent to Wimbledon. 'Slazenger shut down the factory in the UK years ago and moved the equipment to Bataan in the Philippines', said Mark Johnson. 'They still get the felt from Stroud, as it requires a bit more technical expertise. Shipping wool from New Zealand to Stroud and then sending the felt back to the Philippines adds a lot of miles, but they obviously want to use the best wool for the Wimbledon balls.'

DIAGNOSTIC QUESTION

Are supply chain dynamics under control?

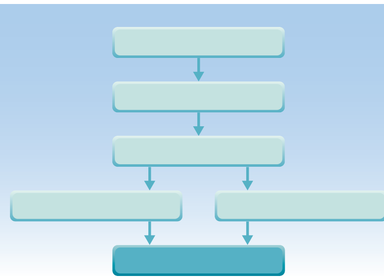

There are dynamics that exist between firms in supply chains that cause errors, inaccuracies and volatility, and these increase for operations further upstream in the supply chain. This effect is known as the bullwhip effect,[11] so called because a small disturbance at one end of the chain causes increasingly large disturbances as it works its way towards the end. Its main cause is a perfectly understandable and rational desire by the different links in the supply chain to manage their levels of activity and inventory sensibly. To demonstrate this, examine the production rate and stock levels for the supply chain shown in Table 7.2. This is a four-stage supply chain where an original equipment manufacturer (OEM) is served by three tiers of suppliers. The demand from the OEM's market has been running at a rate of 100 items per period, but in period 2, demand reduces to 95 items per period. All stages in the supply chain work on the principle that they will keep in stock one period's demand. This is a simplification but not a gross one. Many operations gear their inventory levels to their demand rate. The column headed 'stock' for each level of supply shows the starting stock at the beginning of the period and the finish stock at the end of the period. At the beginning of period 2, the OEM has 100 units in stock (that being the rate of demand up to period 2). Demand in period 2 is 95 and so the OEM knows that it would need to produce sufficient items to finish up at the end of the period with 95 in stock (this being the new demand rate). To do this, it need only manufacture 90 items; these, together with five items taken out of the starting stock, will supply demand and leave a finished stock of 95 items. The beginning of period 3 finds the OEM with 95 items in stock. Demand is also 95 items and therefore its production rate to maintain a stock level of 95 will be 95 items per period. The OEM now operates at a steady rate of producing 95 items per period. Note, however, that a change in demand of only five items has produced a fluctuation of 10 items in the OEM's production rate.

Carrying this same logic through to the first-tier supplier, at the beginning of period 2, the second-tier supplier has 100 items in stock. The demand which it has to supply in period 2 is derived from the production rate of the OEM. This has dropped down to 90 in period 2. The first-tier supplier therefore has to produce sufficient to supply the demand of 90 items (or the equivalent) and leave one month's demand (now 90 items) as its finished stock. A production rate of 80 items per month will achieve this. It will therefore start period 3 with an opening stock of 90 items, but the demand from the OEM has now risen to 95 items. It therefore has

Table 7.2 Fluctuations of production levels along the supply chain in response to small change in end customer demand

	Third-tier supplier		*Second-tier supplier*		*First-tier supplier*		*Original equipment manufacturer*		
Period	*Production*	*Stock*	*Production*	*Stock*	*Production*	*Stock*	*Production*	*Stock*	*Demand*
1	100	100	100	100	100	100	100	100	100
		100		100		100		100	
2	20	100	60	100	80	100	90	100	95
		60		80		90		95	
3	180	60	120	80	100	90	95	95	95
		120		100		95		95	
4	60	120	90	100	95	95	95	95	95
		90		95		95		95	
5	100	90	95	95	95	95	95	95	95
		95		95		95		95	
6	95	95	95	95	95	95	95	95	95
		95		95		95		95	

Starting stock (a) + production (b) = finishing stock (c) + demand, that is production in previous tier down (d): see explanation in text. Note all stages in the supply chain keep one period's inventory, c = d.

to produce sufficient to fulfil this demand of 95 items and leave 95 items in stock. To do this, it must produce 100 items in period 3. After period 3 the first-tier supplier then resumes a steady state, producing 95 items per month. Note again, however, that the fluctuation has been even greater than that in the OEM's production rate, decreasing to 80 items a period, increasing to 100 items a period, and then achieving a steady rate of 95 items a period. Extending the logic back to the third-tier supplier, it is clear that the further back up the supply chain an operation is placed, the more drastic are the fluctuations.

OPERATIONS PRINCIPLE

Demand fluctuations become progressively amplified as their effects work back up the supply chain.

This relatively simple demonstration ignores any time lag in material and information flow between stages. In practice there will be such a lag, and this will make the fluctuations even more marked. Figure 7.9 shows the net result of all these effects in a typical supply chain. Note the increasing volatility further back in the chain.

Controlling supply chain dynamics

The first step in improving supply chain performance involves attempting to reduce the bullwhip effect. This usually means coordinating the activities of the operations in the chain in several ways:[12]

Share information throughout the supply chain

One reason for the bullwhip effect is that each operation in the chain reacts only to the orders placed by its *immediate* customer. They have little overview of what is happening throughout the chain. But if chain-wide information is shared throughout the chain, it is unlikely that such wild fluctuations will occur. With information transmitted throughout the chain, all the

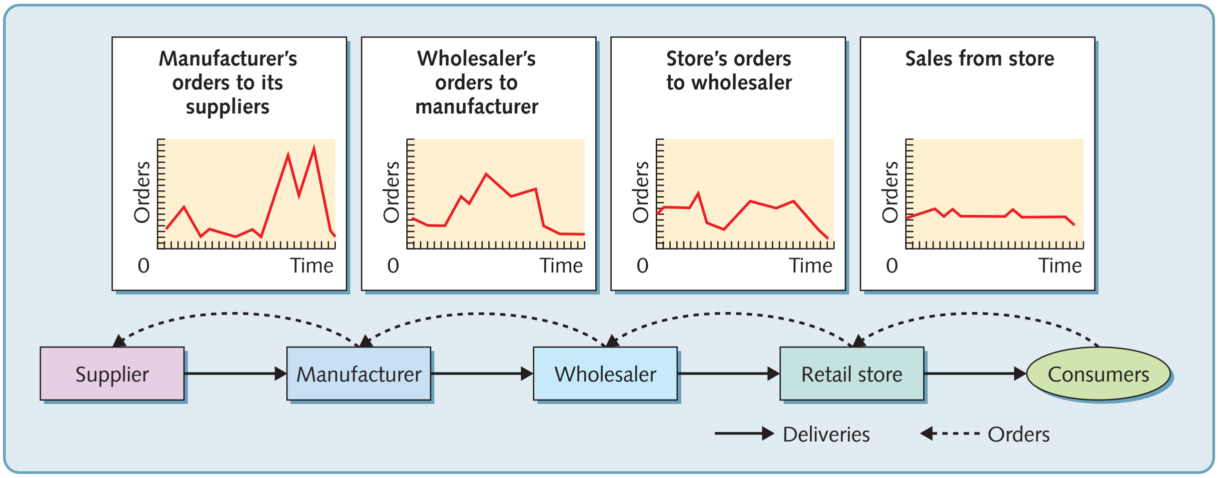

Figure 7.9 Typical supply chain dynamics

operations can monitor true demand, free of distortions. So, for example, information regarding supply problems, or shortages, can be transmitted down the chain so that downstream customers can modify their schedules and sales plans accordingly. For example, the electronic point-of-sale (EPOS) systems, used by many retailers, make information on current demand downstream in the supply chain available to upstream operations. Sales data from checkouts or cash registers are consolidated and transmitted to the warehouses, transportation companies and supplier operations in the supply chain. This means that suppliers can be aware of the 'real' movements in the market.

OPERATIONS PRINCIPLE

The bullwhip effect can be reduced by information sharing, aligning planning and control decisions, improving flow efficiency and better forecasting.

Align all the channels of information and supply

Channel alignment means the adjustment of scheduling, material movements, stock levels, pricing and other sales strategies so as to bring all the operations in the chain into line with each other. This goes beyond the provision of information. It means that the systems and methods of planning and control decision-making are harmonised through the chain. For example, even when using the same information, differences in forecasting methods or purchasing practices can lead to fluctuations in orders between operations in the chain. One way of avoiding this is to allow an upstream supplier to manage the inventories of its downstream customer. This is known as vendor-managed inventory (VMI). So, for example, a packaging supplier could take responsibility for the stocks of packaging materials held by a food manufacturing customer. In turn, the food manufacturer takes responsibility for the stocks of its products that are held in its customer's, the supermarket's warehouses.

Increase operational efficiency throughout the chain

'Operational efficiency' in this context means the efforts that each operation in the chain makes to reduce its own complexity, the cost of doing business with other operations in the chain, and its throughput time. The cumulative effect of this is to simplify throughput in the whole chain. For example, imagine a chain of operations whose performance level is relatively poor: quality defects are frequent, the lead time to order products and services is long, delivery is unreliable, and so on. The behaviour of the chain would be a continual sequence of errors and effort wasted in replanning to compensate for the errors. Poor quality would mean extra and unplanned orders being placed, and unreliable delivery and slow delivery lead times would mean high safety stocks. Just as important, most operations managers' time would be spent coping with the inefficiency. By contrast, a chain whose operations had high levels of operations performance would be more predictable and have faster throughput, both of which would help to minimise supply chain fluctuations.

Improve forecasts

Improved forecast accuracy also helps to reduce the bullwhip effect. Bullwhip is caused by the demand pattern, lead times, forecasting mechanisms and the replenishment decisions used to order product from production facilities or suppliers. Improving the accuracy of your forecasts directly reduces the inventory holding requirements that will achieve customer service level targets. Reducing lead times means that you need to forecast less far into the future and thus lead times have a large impact on bullwhip and inventory costs. The exact nature of how bullwhip propagates in a supply chain is also dependent on the nature of the demand pattern. Negatively correlated demands require less inventory in the supply chain than positively correlated demand patterns, for example. But bullwhip is not unavoidable. By using sophisticated replenishment policies, designed using control engineering principles, many businesses have been able to eliminate bullwhip effects. Sometimes this comes at a cost. Extra inventory may be required in parts of the chain, or customer service levels reduce. But more often bullwhip avoidance creates a 'win-win'. It reduces inventory requirements and improves customer service.

Critical commentary

This emphasis on understanding the end customer in a supply chain has led some authorities to object to the very term *supply* chain. Rather, they say, they should be referred to as *demand* chains. Their argument is based on the idea that the concept of 'supply' implies a 'push' mentality. Any emphasis on pushing goods through a supply chain should be avoided. It implies that customers should consume what suppliers see fit to produce. On the other hand, referring to 'demand chains' puts proper emphasis on the importance of seeing customers as pulling demand through the chain. Nevertheless, 'supply chain' is still the most commonly used term.

- The use of technology in supply chain management is not always universally welcomed. Even e-procurement is seen by some as preventing closer partnership-type relationships that, in the long run, may be more beneficial. Similarly, track-and-trace technology is seen by some as a waste of time and money. '*What we need,*' they argue, '*is to know that we can trust the delivery to arrive on time, we do not need the capability to waste our time finding out where the delivery is.*'

- The Internet of Things also opens up many ethical issues. People see its potential and its dangers in very different ways. Take the following statement:[13]

'Supermarket cards and other retail surveillance devices are merely the opening volley of the marketers' war against consumers. If consumers fail to oppose these practices now our long term prospects may look like something from a dystopian science fiction novel. . . . though many . . . proponents appear focused on inventory and supply chain efficiency, others are developing financial and consumer applications that, if adopted, will have chilling effects on consumers' ability to escape the oppressive surveillance of

manufacturers, retailers, and marketers. Of course, government and law enforcement will be quick to use the technology to keep tabs on citizens as well.'

- Possibly the IoT issue to make the most headlines is that of its potential for being hacked to launch security attacks. As some point out, connecting things to each other through the internet will always expose new vulnerabilities. Looked at one way, we no longer have things with computers embedded in them. We have computers with things attached to them. Security is, arguably, given too little emphasis in the use of IoT in supply chain operations.[14]

SUMMARY CHECKLIST

- ☐ Is it understood that the performance of any one operation is partly a function of all the other operations in the supply chain?
- ☐ Are supply chain concepts applied internally as well as externally?
- ☐ Are supply chain objectives understood in the context of the whole chain rather than the single operation?
- ☐ Which product or service groups are 'functional' and which are 'innovative'?
- ☐ Therefore, which products or service groups need 'lean' and which need 'agile' supply chain management?
- ☐ Is the position on the' transactional to partnership' spectrum understood for each customer and supplier relationship?
- ☐ Are customer and supplier relationships at an appropriate point on the transactional to partnership spectrum?
- ☐ Are 'partnership' relationships *really* partnerships or are they just called that?
- ☐ Are suppliers and potential suppliers rigorously assessed using some scoring procedure?
- ☐ Are the trade-offs inherent in supplier selection understood?
- ☐ Is the approach to single or multi-sourcing appropriate?
- ☐ Is the purchasing activity making full use of internet-based mechanisms?
- ☐ Are service-level agreements used? And do they develop over time?
- ☐ Is sufficient effort put into supplier development?
- ☐ Are actual and potential mismatches of perception in the supplier relationships explored?
- ☐ Is the difference between dependent and independent demand understood?
- ☐ Is the potential for outsourcing logistics services regularly explored?
- ☐ Could the adoption of new technologies such as IoT have any benefit?
- ☐ Has the idea of customer development been explored?
- ☐ Have mechanisms for reducing the impact of the bullwhip effect been explored?
- ☐ Has there been a risk assessment to assess supply chain vulnerability?

CASE STUDY

Supplying fast fashion[15]

Introduction

Garment retailing has changed. No longer is there a standard look that all retailers adhere to for a whole season. Fashion is fast, complex and furious. Different trends overlap and fashion ideas that are not even on a store's radar screen can become 'must haves' within six months. Many retail businesses with their own brands such as H&M and Zara, sell up-to-the-minute fashion and low prices, in stores that are clearly focused on one particular market. In the world of fast fashion, catwalk designs speed their way into high street stores at prices anyone can afford. The quality of the garment means that it may only last one season, but fast fashion customers don't want yesterday's trends. As *Newsweek* puts it, '. . . .*.being a 'quicker picker-upper' is what made fashion retailers H&M and Zara successful. [They] thrive by practicing the new science of 'fast fashion'; compressing product development cycles as much as six times.'* But the retail operations that customers see are only the end part of the supply chains that feed them. And these have also changed.

At its simplest level, the fast fashion supply chain has four stages. First, the garments are designed, after which they are manufactured. They are then distributed to the retail outlets where they are displayed and sold in retail operations designed to reflect the business's brand values. In this short case, we examine two fast fashion operations, Hennes and Mauritz (known as H&M) and Zara, together with United Colours of Benetton (UCB), a similar chain but with a different market positioning.

Benetton

Almost 50 years ago Luciano Benetton took the world of fashion by storm by selling the bright, casual sweaters designed by his sister across Europe (and later the rest of the world), promoted by controversial advertising. By 2005 the Benetton Group was present in 120 countries throughout the world. Selling casual garments, mainly under its UCB and its more fashion-orientated Sisley brands, it produces 110 million garments a year, over 90 per cent of them in Europe. Its retail network of over 5,000 stores produces revenue of around €2 billion. Benetton products are seen as less 'high fashion' but higher quality and durability, with higher prices, than H&M and Zara.

H&M

Established in Sweden in 1947, H&M now sell clothes and cosmetics in over 1,000 stores in 20 countries around the world. The business concept is 'fashion and quality at the best price'. With more than 40,000 employees, and revenues of around SEK 60,000 million, its biggest market is Germany, followed by Sweden and the UK. H&M are seen by many as the originator of the fast fashion concept. Certainly, they have years of experience at driving down the price of up-to-the-minute fashions. '*We ensure the best price,'* they say, *'by having few middlemen, buying large volumes, having extensive experience of the clothing industry, having a great knowledge of which goods should be bought from which markets, having efficient distribution systems, and being cost-conscious at every stage.'*

Zara

The first store opened almost by accident in 1975 when Amancio Ortega Gaona, a women's pyjama manufacturer, was left with a large cancelled order. The shop he opened was intended only as an outlet for cancelled orders. Now, Inditex, the holding group that includes the Zara brand, has over 1,300 stores in 39 countries with sales of over €3 billion. The Zara brand accounts for over 75 per cent of the group's total retail sales, and is still based in North West Spain. By 2003 it had become the world's fastest growing volume garment retailer. The Inditex group also has several other branded chains including Pull and Bear, and Massimo Dutti. In total, it employs almost 40,000 people in a business that is known for a high degree of vertical integration compared with most fast fashion companies. The company believes that it is their integration along the supply chain that allows them to respond to customer demand fast and flexibly, while keeping stock to a minimum.

Design

All three businesses emphasise the importance of design in this market. Although not *haute couture,* capturing design trends is vital to success. Even the boundary between high and fast fashion is starting to blur. In 2004 H&M recruited high-fashion designer Karl Lagerfeld, previously noted for his work with more exclusive brands. For H&M, his designs were priced for value rather than exclusivity, *'Why do I work for H&M? Because I believe in inexpensive clothes, not 'cheap' clothes,'* said Lagerfeld. Yet most of H&M's products come from over a hundred designers in Stockholm who work with a team of 50 pattern designers, around 100 buyers and a number of budget controllers. The department's task is to find the optimum balance between the three components comprising H&M's business concept: fashion, price and quality. Buying volumes and delivery dates are then decided.

Zara's design functions are organised in a different way to most similar companies. Conventionally, the design input comes from three *separate* functions: the designers themselves, market specialists, and buyers who place orders on to suppliers. At Zara the design stage is split into three product areas: women's, men's and children's garments. In each area, designers, market specialists and buyers are co-located in design halls that also contain small workshops for trying out prototype designs. The market specialists in all three design halls are in regular contact with Zara retail stores, discussing customer reaction to new designs. In this way, the retail stores are not the end of the whole supply chain but the beginning of the design stage of the chain. Zara's designers, around 300 in total with an average age of 26, produce approximately 40,000 items per year, of which about 10,000 go into production.

Benetton also has about 300 designers, who not only design for all their brands, but also are engaged in researching new materials and clothing concepts. Since 2000 the company has moved to standardise their range globally. At one time more than 20 per cent of its ranges were customised to the specific needs of each country, now only between 5 and 10 per cent of garments are customised. This reduced the number of individual designs offered globally by over 30 per cent, strengthening the global brand image and reducing production costs.

Both H&M and Zara have moved away from the traditional industry practice of offering two 'collections' a year, for Spring/Summer and Autumn/Winter. Their 'seasonless cycle' involves the continual introduction of new products on a rolling basis throughout the year. This allows designers to learn from customers' reactions to their new products and incorporate them quickly into more new products. The most extreme version of this idea is practised by Zara. A garment will be designed; a batch manufactured and 'pulsed' through the supply chain. Often the design is never repeated, it may be modified and another batch produced, but there are no 'continuing' designs, as such. Even Benetton, have increased the proportion of what they call 'flash' collections, small collections that are put into its stores during the season.

Manufacturing

At one time Benetton focused its production on its Italian plants. Then it significantly increased its production outside Italy to take advantage of lower labour costs. Non-Italian operations include factories in North Africa, Eastern Europe and Asia. Yet each location operates in a very similar manner. A central, Benetton-owned operation performs some manufacturing operations (especially those requiring expensive technology) and coordinates the more labour-intensive production activities that are performed by a network of smaller contractors (often owned and managed by ex-Benetton employees). These contractors may in turn subcontract some of their activities. The company's central facility in Italy allocates production to each of the non-Italian networks, deciding what and how much each is to produce. There is some specialisation, for example, jackets are made in Eastern Europe while T-shirts are made in Spain. Benetton also has a controlling share in its main supplier of raw materials, to ensure fast supply to its factories. Benetton are also known for the practice of dying garments after assembly rather than using dyed thread or fabric. This postpones decisions about colours until late in the supply process so that there is a greater chance of producing what is needed by the market.

H&M does not have any factories of its own, but instead works with around 750 suppliers. Around half of production takes place in Europe and the rest mainly in Asia. It has 21 production offices around the world that between them are responsible for coordinating the suppliers who produce over half a billion items a year for H&M. The relationship between production offices and suppliers is vital, because it allows fabrics to be bought in early. The actual dyeing and cutting of the garments can then be decided at a later stage in the production The later an order can be placed on suppliers, the less the risk of buying the wrong thing. Average supply lead times vary from 3 weeks up to 6 months, depending on the nature of the goods. However, *'The most important thing,'* they say, *'is to find the optimal time to order each item. Short lead times are not always best. Some high-volume fashion basics, it is to our advantage to place orders far in advance. Trendier garments require considerably shorter lead times.'*

Zara's lead times are said to be the fastest in the industry, with a 'catwalk to rack' time as little as 15 days. According to one analyst this is because they *'owned most of the manufacturing capability used to make their products, which they use as a means of exciting and* *stimulating customer demand.'* About half of Zara's products are produced in its network of 20 Spanish factories, which, like at Benetton, tended to concentrate on the more capital intensive operations such as cutting and dyeing. Subcontractors are used for most labour-intensive operations like sewing. Zara buys around 40 per cent of their fabric from its own wholly-owned subsidiary, most of which is in undyed form for dying after assembly. Most Zara factories and their subcontractors work on a single shift system to retain some volume flexibility.

Distribution

Both Benetton and Zara have invested in highly automated warehouses, close to their main production centres that store, pack and assemble individual orders for their retail networks. These automated warehouses represent a major investment for both companies. In 2001, Zara caused some press comment by announcing that it would open a second automated warehouse even though, by its own calculations, it was only using about half its existing warehouse capacity. More recently, Benetton caused some controversy by announcing that it was exploring the use of RFID tags to track its garments.

At H&M, while the stock management is primarily handled internally, physical distribution is subcontracted. A large part of the flow of goods is routed from production site to the retail country via H&M's transit terminal in Hamburg. Upon arrival the goods are inspected and allocated to the stores or to the centralised store stock room. The centralised store stock room, within H&M referred to as 'Call-Off warehouse' replenishes stores on item level according to what is selling.

Retail

All H&M stores (average size, 1,300 square metres) are owned and solely run by H&M. The aim is to '*create a comfortable and inspiring atmosphere in the store that makes it simple for customers to find what they want and to feel at home*'. This is similar to Zara stores, although they tend to be smaller (average size, 800 square metres). Perhaps the most remarkable characteristic of Zara stores is that garments rarely stay in the store for longer than 2 weeks. Because product designs are often not repeated and produced in relatively small batches, the range of garments displayed in the store can change radically every 2 or 3 weeks. This encourages customers both to avoid delaying a purchase and to revisit the store frequently.

Since 2000 Benetton have been reshaping their retail operations. At one time the vast majority of Benetton retail outlets were small shops run by third parties. Now these small stores have been joined by several, Benetton-owned and operated, larger stores (1,500 to 3,000 square metres). These mega-stores can display the whole range of Benetton products and reinforce the Benetton shopping experience.

Question

Compare and contrast the approaches taken by H&M, Benetton and Zara to managing their supply chain.

APPLYING THE PRINCIPLES

Some of these exercises can be answered by reading the chapter. Others will require some general knowledge of business activity and some might require an element of investigation. **Model answers for the first two questions are available on the student companion website.** For model answers to the other questions in this section, please ask your tutor/lecturer.

1 The COO of Super Cycles was considering her sourcing strategy. *'I have two key questions, for each of our outsourced parts, what is the risk in the supply market, and what is the criticality of the product or service to our business? As far as risk is concerned, we can consider the number of alternative suppliers, how easy it is to switch from one supplier to another, possible exit barriers, and the cost of bringing operations back in-house. As far as criticality is concerned, we can consider each component's importance in terms of volume purchased, percentage of total purchase cost, and the impact on business growth.'* Four key outsourced components are shown in Table 7.3. What approach to sourcing these components would you recommend?

Table 7.3 Four outsourced components for Super Cycles

Component	*Cost (as a proportion of total material cost)*	*Suppliers*	*Ease of changing supplier*
The innertubes	3%	Many alternative suppliers	Very easy. Could do it in days
Frame tubing	15%	Only 1 supplier capable at the moment. Could take a long time to develop a new supplier	Difficult in the short term, possible in the longer term
Carbon fibre stem and bars	32%	Relatively large number of available suppliers	Relatively easy. Would probably take a few weeks for new contract
'Groupset' gearing system	35%	Few suppliers who are capable of manufacturing these components to sufficient quality	Complex to source. Could swich supplier in the longer term, but would pose quality risk

2 A chain of women's apparel retailers had all their products made by Lopez Industries, a small but high-quality garment manufacturer. They worked on the basis of two seasons: Spring/Summer season and Autumn/Winter. '*Sometimes we are left with surplus items because our designers have just got it wrong*', said the retailer's chief designer. '*It is important that we are able to flex our order quantities from Lopez during the season. Although they are a great supplier in many ways, they can't change their production plans at short notice.*' Lopez Industries was aware of this. '*I know that they are happy with our ability to make even the most complex designs to high level of quality. I also know that they would like us to be more flexible in changing our volumes and delivery schedules. I admit that we could be more flexible within the season. Partly, we can't do this because we have to buy in cloth at the beginning of the season based on the forecast volumes from our customers. Even if we could change our production schedules, we could not get extra deliveries of cloth. We only deal with high-quality and innovative cloth manufacturers who are very large compared to us, so we do not represent much business for them.*' A typical cloth supplier said: '*we compete primarily on quality and innovation. Designing cloth is as much of a fashion business as designing the clothes into which it is made. Our cloth goes to tens of thousands of customers around the world. These vary considerably in their requirements, but presumably all of them value our quality and innovation.*' How should the retailer try to influence this supply chain in order to improve its performance?

3 The example of the bullwhip effect shown in Table 7.2 shows how a simple 5 per cent reduction in demand at the end of the supply chain causes fluctuations that increase in severity the further back an operations is placed in the chain.

(a) Using the same logic and the same rules (i.e. all operations keep one period's inventory), what would the effect on the chain be if demand fluctuated period by period between 100 and 95? That is, period 1 has a demand of 100, period 2 has a demand of 95, period 3 a demand of 100, period 4 a demand of 95, and so on?

(b) What happens if all operations in the supply chain decided to keep only half of the periods demand as inventory?

(c) Find examples of how supply chains try to reduce this bullwhip effect.

4 If you were the owner of a small local retail shop, what criteria would you use to select suppliers for the goods that you wish to stock in your shop? Visit three shops that are local to you and ask the owners how they select their suppliers. In what way were their answers different from what you thought they might be?

5 Many companies devise a policy on ethical sourcing covering such things as workplace standards and business practices, health and safety conditions, human rights, legal systems, child labour, disciplinary practices, wages and benefits, etc. (a) What do you think motivates a company to draw up a policy of this type? (b) What other issues would you include in such a supplier selection policy?

Notes on chapter

1 Sources include: Butler, S. (2017) 'Hand delivered: will Ocado's robot soon be picking your shopping?', the *Guardian,* 31 January; Farrell, S. (2014) 'Ocado dismisses fears over Waitrose supply deal', the *Guardian,* 11 September: Pratley, N. (2012) 'Ocado: Buy two problems get one free', the *Guardian,* 23 January.

2 Sources include: Bajarin, T. (2016) 'Apple's little-known (and somewhat unsexy) secret to success', recode.nat, posted 30 August; Williams, R. (2014) 'Apple's supply chain in trouble?', *The Telegraph,* 23 March; Ellinor, R. (2013) 'Apple named world's best supply chain', supplymamagement.com, 24 May; Tate, R. (2014) 'Apple tightens grip on worldwide supply chain with Asia hiring spree', *Wired Magazine,* 3 March.

3 The UK Government Department for Environment, Food and Rural Affairs (2006) 'Procuring the Future – report of the Sustainable Procurement Task Force'.

4 The North Face Website, http://uk.thenorthface.com/blog/uk/en/sustainability

5 Fisher, M.L. (1997) 'What is the right supply chain for your product', *Harvard Business Review,* March–April.

6 Kapour, V. and Gupta, A. (1997) 'Aggressive sourcing: a free-market approach, *Sloan Management Review,* Fall.

7 Parkhe, A. (1993) 'Strategic alliance structuring', *Academy of Management Journal,* vol. 36, 794–829.

8 Grad, C. (2000) 'A network of supplies to be woven into the web', *Financial Times,* 9 February.

9 Lee, L. and Dobler, D.W. (1977) *Purchasing and Materials Management,* McGraw-Hill.

10 Harland, C.M. (1996) 'Supply chain management relationships, chains and networks', *British Journal of Management*, vol. 1, no. 7.

11 Lee, H.L., Padmanabhan, V. and Whang, S. (1997) 'The bullwhip effect in supply chains', *Sloan Management Review*, Spring.

12 Thanks to Stephen Disney at Cardiff Business School, UK, for help with this section.

13 MIT Auto-ID website, and Albercht, K. (2002) 'Supermarket cards: tip of the surveillance iceberg', *Denver University Law Review*, June.

14 See for example, the views of Bruce Schneier (2017) 'Security and the internet of things', Schneier on Scecurity Blog, https://www.schneier.com/blog/

15 All data from public sources and reflect period 2004–2005.

TAKING IT FURTHER

Akkermans, H. and Voss, C. (2013) '*The service bullwhip effect*', *International Journal of Operations & Production Management*, *vol. 33, no. 6, 765–788.* *An academic paper that deals with the important issue of service supply.*

Bolstorff, P. and Rosenbaum, R. (2008) *Supply Chain Excellence – a handbook for dramatic improvement using the SCOR model (2nd edn), American Management Association.*

Bolstorff, P. (2004) '*Supply Chain by the Numbers*', *Logistics Today, July, 46–50.*

Chopra, S. and Meindl, P. (2015) *Supply Chain Management: Strategy, Planning and Operations, Pearson.* *A good textbook that covers both strategic and operations issues.*

Christopher, M. (2016) *Logistics and Supply Chain Management (5th edn)*, Financial Times Prentice Hall. *A comprehensive treatment on supply chain management from a distribution perspective by one of the gurus of supply chain management.*

Fisher, M.L. (1997) '*What is the Right Supply Chain for your Product?*', *Harvard Business Review, vol. 75, no. 2.* *A particularly influential article that explores the issue of how supply chains are not all the same.*

Parmigiania A, Klassenb, R.D. and Russoa, M.V. (2011) '*Efficiency meets accountability: Performance implications of supply chain configuration, control, and capabilities*', *Journal of Operations Management*, *vol. 29, 212–223.* *Excellent review and discussion of broad supply issues.*

Presutti Jr., William, D. and Mawhinney, J. R. (2007) '*The Supply Chain Finance Link*', *Supply Chain Management Review, September, 32–38.*

8 Capacity management

Introduction

At its core, the role of capacity management is to reconcile the demand and supply for products and services. To do this, managers must understand and cope with two competing requirements. One is to understand the importance of delivering products and services to customers quickly and reliably. This maintains customer satisfaction, and (usually) has a positive effect on revenue. The other, is the need for operations and their extended supply networks to maintain efficiency by minimising excess capacity costs. And this is why capacity management is so important – it has an impact on both revenue and costs, and therefore profitability (or the general effectiveness of service delivery in not-for-profit operations). In this chapter, we look primarily at these competing tensions at an aggregated and medium-term level. Figure 8.1 shows the position of the ideas described in this chapter in the general model of operations management.

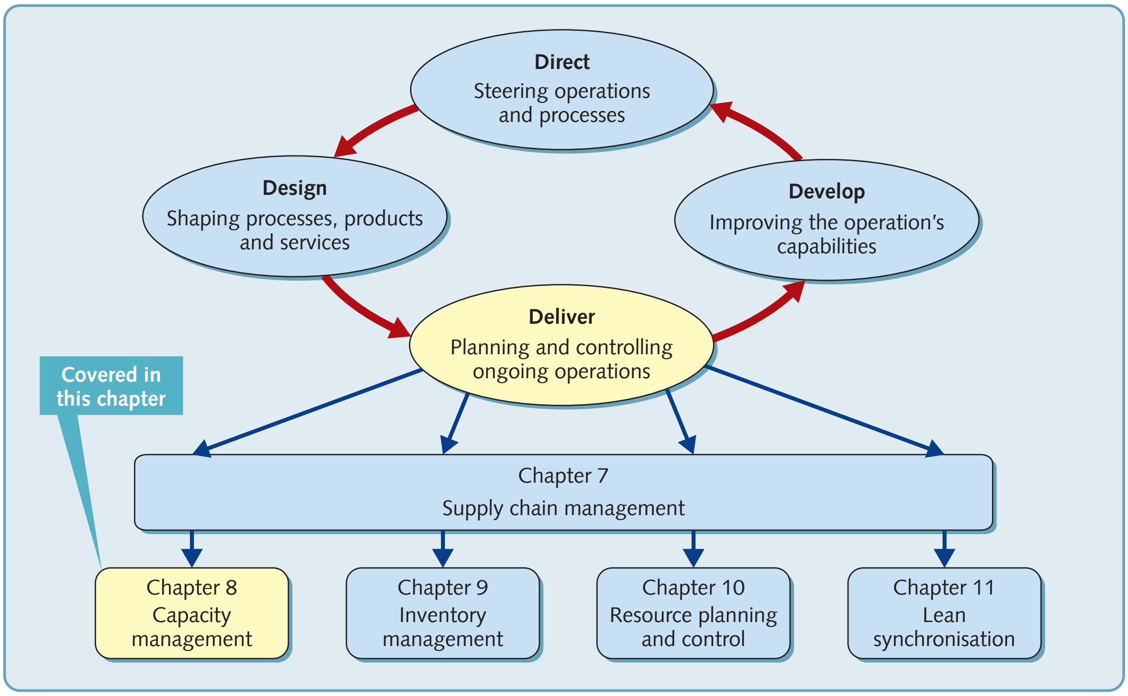

Figure 8.1 Capacity management is the activity of coping with mismatches between demand and the ability to supply demand

EXECUTIVE SUMMARY

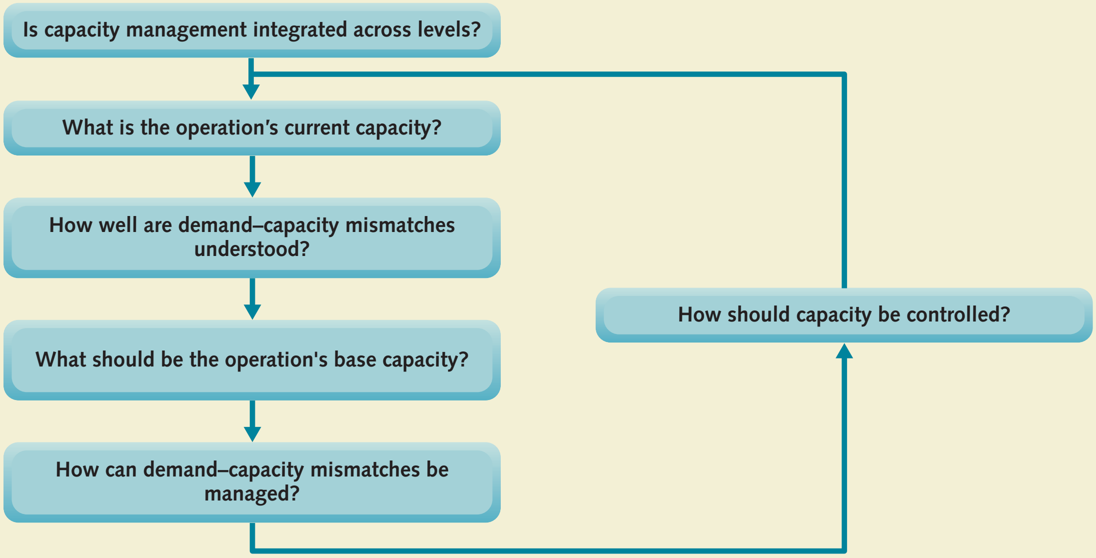

Is capacity management integrated across levels?

Capacity management is the activity of understanding the nature of demand and supply, and of selecting alternative supply side (level- and chase-capacity plans) and demand side (demand management and yield management) responses to effectively meet the needs of customers, while maintaining resource efficiency. Every operation and process needs to know their capacity because if they have too little they cannot meet demand (bad for revenue), and if they have too much they are paying for more capacity than they need (bad for costs). However, capacity needs to be managed at different levels – long, medium and short term. In the long term, capacity decisions include total capacity, its distribution between sites, and its location. In the medium term, capacity decisions include the extent to which capacity levels fluctuate, staffing levels, and the degree of subcontracting. In the short term, capacity decisions include which resources should be allocated to each set tasks, and how tasks should be allocated to resources. Although capacity decisions are taken for different timescales and spanning different areas of the operation, each level of capacity decision is made with the constraints of a higher level. Moreover, the activity of managing capacity should provide feedback to the higher level.

What is the operation's current capacity?

The first step in managing capacity is to be able to measure it, which can be relatively difficult unless the operation is standardised and repetitive. Any measure of capacity will contain a number of assumptions, each of which may be necessary to give an estimate of capacity, but each of which obscures some aspect of reality. One assumption relates to the mix of products or services supplied, that is, what it is being required to do. Another is the time over which they are supplied. The level of activity and output that may be achievable over short periods of time is not the same as the capacity that is sustainable on a regular basis. It also depends on the specification of what is supplied. Some services can increase their output by changing the specification of their output. The reduction in capacity is sometimes called 'capacity leakage'; one method of assessing this leakage is called the overall equipment effectiveness (OEE) measure.

How well are demand–capacity mismatches understood?

Understanding the nature of potential demand–capacity mismatches is central to capacity management. A key issue is the nature of demand and capacity fluctuations, especially the degree to which they are predictable. If fluctuations are predictable, they can be planned in advance to minimise their costs. If fluctuations are unpredictable, the main objective is to react to them quickly. Accurate simple forecasting is an advantage because it converts unpredictable variation into predictable variation. However, a broader approach to enhancing market knowledge generally can reveal more about the options for managing mismatches.

What should be the operation's base capacity?

Capacity planning often involves setting a base level of capacity and then planning capacity fluctuations around it. The level at which base capacity is set depends on three main factors: the relative importance of the operation's performance objectives; the perishability of the operation's outputs; and the degree of variability in demand or supply. High service levels, high perishability of an operation's outputs and a high degree of variability either in demand or supply, all indicate a relatively high level of base capacity.

How can demand–capacity mismatches be managed?

Demand–capacity mismatches usually call for some degree of capacity adjustment over time. There are three pure methods of achieving this, although in practice a mixture of all three may be used. A 'level capacity' plan involves no change in capacity and requires that the operation absorb demand–capacity mismatches, usually through under- or over-utilisation of its resources, or the use of inventory. The 'chase demand' plan involves the changing of capacity through such methods as overtime, varying the size of the work force, subcontracting, and so on. The 'manage demand' plan involves an attempt to change demand through pricing or promotion methods, or changing product or service mix to reduce fluctuations in activity levels. When outputs cannot be stored, yield management is a common method of coping with mismatches. When outputs can be stored, cumulative representations are a useful method for planning capacity.

How should capacity be controlled?

In practice, capacity management is a dynamic process with decisions reviewed period by period. It is essential that capacity decisions made in one period reflect the knowledge accumulated from experiences in previous periods.

DIAGNOSTIC QUESTION

Is capacity management integrated across levels?

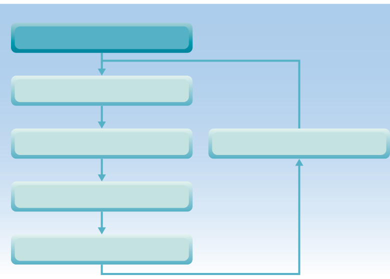

The capacity of an operation (or single process) is the output that it can deliver in a defined unit of time. It reflects an operation's 'ability to supply', at least in a quantitative sense. Capacity management is the activity of understanding the nature of demand and supply (capacity), and the mismatches between them. Demand is the quantity of products or services that customers request from an operation or process at any point in time. A mismatch between demand and capacity can occur because demand fluctuates over time, or capacity fluctuates over time, or both.

Defining capacity as 'the ability to supply' is taking a broad view of the term. The 'ability to supply' depends not only on the limitations of the previous stage in a supply network, operation, or process, but also on all the stages up to that point. So, for example, the capacity of an ice cream manufacturer is a function not only of how much ice cream its factories can produce at any point in time, but also of how much packaging material, raw material supplies, and so on, that its suppliers can provide. It may have the factories to make 10,000 kilos of ice cream a day, but if its suppliers of dairy produce can only supply 7,000 kilos a day, then the effective capacity (in terms of 'ability to supply') is only 7,000 kilos per day. Of course, if demand remains steady any operation will attempt to make sure that its supply capacity does not limit its own ability to supply. However, given that capacity management is concerned with fluctuations in demand *and* supply, the process of managing capacity is essentially a dynamic issue.

OPERATIONS PRINCIPLE

Any measure of capacity should reflect the ability of an operation or process to supply demand.

It is worth noting that 'coping' with mismatches between demand and capacity does not always mean that capacity should match demand. Many operations take the deliberate decision to fail to meet demand or fail to fully exploit its ability to supply. For example, a hotel may not make any effort to meet demand in peak periods because doing so would incur unwarranted capital costs. It is therefore content to leave some demand unsatisfied, although it may increase its prices to reflect this. Similarly, a flower grower may not supply the entirety of its potential supply (crop) if doing so would simply depress market prices and reduce its total revenue.

Integrating across the levels of capacity management

The central issue of capacity – deciding if and how to meet demand – can be considered at different organisational levels and over different time periods. Capacity decisions are taken minute-by-minute, day-by-day, month-on-month and year-on-year. They are also made within the constraints of the physical limits of the operation, the ability of its suppliers to supply, the availability of staff, and so on. In other words, although capacity decisions are taken for different timescales and spanning different areas of the operation, each level of capacity decision is made with the constraints of a higher level. Moreover, the activity of managing capacity should provide feedback to the higher level. Figure 8.2 summarises these different levels of capacity decision-making at three levels.

At a strategic level, capacity decisions include (see Chapter 4):

- How much capacity do we need in total?
- How should the capacity be distributed?
- Where should the capacity be located?

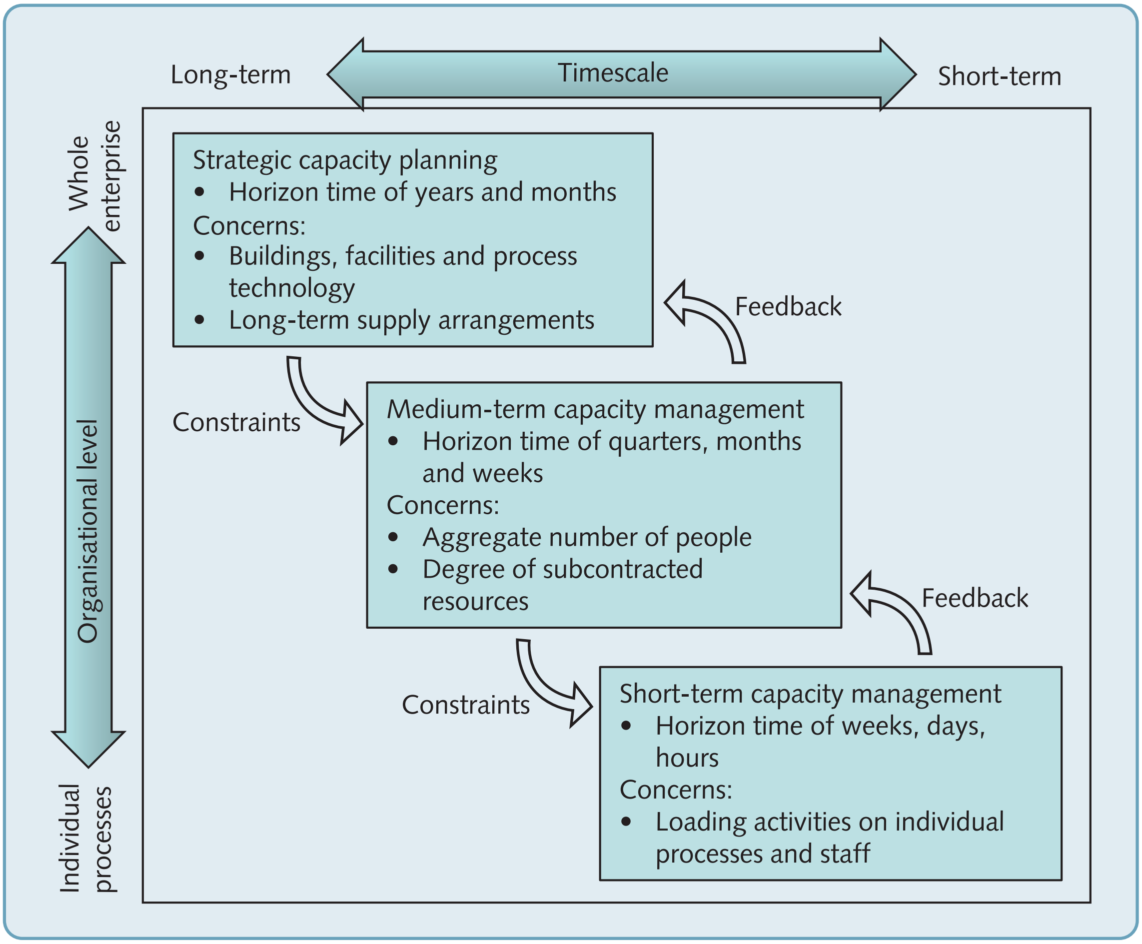

Figure 8.2 Capacity management should be integrated across levels because each level constrains what can be done in the level below and can provide feedback for the level above

At a medium-term level, capacity decisions include:

- To what extent do we keep capacity level or fluctuate capacity levels?
- Should we change staffing levels as demand changes?
- Should we subcontract or off-load demand?

At a short-term level, capacity decisions include (see Chapter 10):

- Which resources should be allocated to each set tasks?
- When should tasks be allocated to resources?

However, these three levels of capacity decision-making are, to some extent, arbitrary. There will always be overlap between the levels, and the actual timescales of the three levels will vary between industries. The important point here is that capacity planning should be integrated across levels because each level constrains what can be done in the level below.

The following two examples illustrate the nature of capacity management.

EXAMPLE

Christmas at Amazon[1]

Amazon is probably the best-known online retailer in the world. It owes much of its success to the network of highly efficient fulfilment operations based across Europe, North America, Latin America and Asia. It is these centres that organise the shipment of millions of individual items and enable such an appealing offer of product range, low price and quick delivery. Typical of Amazon's

shipment operations is its 46,000 m^2 Marston Gate warehouse in Milton Keynes, one of eight in the UK. In the warehouse, products are stocked within its extensive shelving arrangement with the position of every item referenced using a portable satellite navigation system. Amazon says that it saves time when their staff retrieve items. '*The computer shows the shortest pick-path*', said Arthur Valdez, the vice-president of Amazon's British operations. The item is then scanned and picked, after which it moves along a conveyor belt to be packed or gift wrapped and then labelled. At this point, an email is automatically sent to the customer informing them that their product is on its way.

Mr Valdez manages a network of fast-moving operations that must always maintain a tight control of their activities; but no time is more testing than the run-up to Christmas. The gift-buying habits of western consumers mean that up to 40 per cent of annual sales value can be made in the final three months of the year, with half of the multibillion online Christmas sales taking place during the end of November and the first two weeks of December. The number of articles sold each day soars from an average of 300,000 to 3.6 million at its peak. In the UK, this day – called 'Cyber Monday' – is at the beginning of December, to be more precise, 9 p.m. on that day, when shoppers, who have normally been paid for the month and spent the weekend browsing the high street, return from work to begin their Christmas shopping in earnest. It makes for a hectic time, '*A full truck is dispatched every three minutes and 24 seconds on our busiest trading day*', said Mr Valdez. But careful forecasting can at least stop the Christmas peak being a surprise. And careful monitoring of customer behaviour has revealed a further trend – after 'Cyber Monday', now comes 'Boomerang Thursday', when customers start to return their unwanted items. '*As the online retail sector continues to grow, so too has consumer demand and confidence to return items, often before Christmas*', said Mark Lewis, chief executive of CollectPlus, which allows customers to return items to a local convenience store. '*This suits retailers. They want to get* [items] *back as soon as possible, so they can sell them on.*' Mark Lewis said that half of his customers return items at off-peak times. '*It peaks at 7 p.m. It reflects how we live our lives these days.*'

However, some retail analysts believe that the advance of technology in the form of mobile phone transactions and broadband has also meant that the significance of 'Cyber Monday' and 'Boomerang Thursday' will diminish because such technology makes it easier to stagger transactions. But for Mr Valdez, it is continual vigilance that allows Amazon to keep up with demand trends. '*Every year it feels like [Christmas starts on] January 1. We are all-year-long focused on understanding the lessons learnt from the previous Christmas,*' he said.

EXAMPLE

Panettone: how Italy's bakers cope with seasonal demand[2]

Panettone has become a national symbol of the Italian Christmas. The light and fluffy, dome-shaped, confection, is dotted with sultanas and candied citrus peel, and is *the* Italian Christmas cake. Traditionally made in Milan, Italy, about 40 million of them are consumed throughout Italy over the holiday period. Now, they are becoming popular around the world. Over a million are exported to the US, while an endorsement from Delia Smith, a celebrity chef, caused a surge in demand in Britain with a well-publicised recipe for trifle made with panettone. This boost to production is good news for the big Italian manufacturers, but although volumes are higher, the product is still seasonal, which poses a problem for even the experienced Milanese confectioners. Smaller, 'artisan' producers simply squeeze a few batches of panettone into their normal baking schedules as Christmas approaches. But for the large industrial producers who need to make millions for the Christmas season, it is not possible. And no panettone manufacturer is larger than the Bauli group. It is one of the foremost manufacturers of confectionery in Europe. Founded over 70 years ago and, in spite of its mass production approach, it has a reputation for quality and technological improvement. The company's output of panettone accounts for 38 per cent of Italian sales. The key to its success, according to the company, is in having '*combined the skill of homemade recipes with high technology [and] quality guaranteed by high standards that are unattainable in craftsman production, but that can only be reached by selecting top quality raw materials, by thousands of tests and checks on the entire production line and the production process*'. In fact, the company says that its size is an advantage. '*High investment in research and technology allow us to manage natural fermentation and guarantee a uniform quality that artisanal bakeries find hard to achieve.*'

Although Bauli has diversified into year-round products like croissants and biscuits, it has acquired a leadership role in the production of products for festive occasions. Seasonal cakes account for over 50 per cent of its turnover of around €420 million. And so successful has it been in its chosen markets that in 2009 it bought Motta and Alemagna, the two big Milanese brands that pioneered the manufacture of panettone. So how does Bauli cope with such seasonality? Partly it is by hiring large numbers of temporary seasonal workers to staff its dedicated production lines. At peak times there can be 1,200 seasonal workers in the factory, more than its permanent staff of around 800. It also starts to build up inventories before demand begins to increase for the Christmas peak. Production of panettone lasts about four months, starting in September. '*Attention to ingredients and the use of new technologies in production give a shelf-life of five months without preservatives*', said Michele Bauli, deputy chairman who is a descendent of the firm's founding family. Temporary workers are also hired to bake other seasonal cakes such as the *colomba*, a dove-shaped Easter treat, which keeps them occupied for a month and a half in the spring.

What do they have in common?

The obvious similarity between both these operations is that they have to cope with relatively high levels of fluctuation in demand for their products and services throughout the year. While the variation in demand of course matters hugely from a capacity perspective, what is perhaps more interesting is the *balance* between the predictable and unpredictable elements of this variation. In both case, a large proportion of the variation in demand can be accurately forecast in advance. This allows planned change in capacity, often through short-term seasonal workers, outsourced additional resources, or running processes for longer periods of the day in the run-up to Christmas. What is more challenging for these operations is the *unpredictable* variation in demand. In the case of Amazon, demand spikes for different products can occur due to the marketing campaigns of companies that may sell products on this site, but who are unlikely to inform Amazon of their promotion activities. Likewise, for panettone producers, while the traditional Italian market is predictable, the growing international market is far less so (for example the surge caused by short-term publicity in the UK).

Variation can come from demand, supply, or both

In both these examples, any mismatches between demand and capacity derive from predictable and unpredictable variation in *demand.* However, some operations also have to cope with predictable and unpredictable variation in *capacity* (if it is defined as 'the ability to supply'). For example, Figure 8.3 shows the demand and capacity variation of two businesses. The first is a domestic appliance repair service. Both demand and capacity vary month-on-month. Capacity varies because the field service operatives in the business prefer to take their vacations at particular times of the year. Nevertheless, capacity is relatively stable throughout the year. Demand, by contrast, fluctuates more significantly. It would appear that there are two peaks of demand through the year, with peak demand being approximately twice the level of the low point in

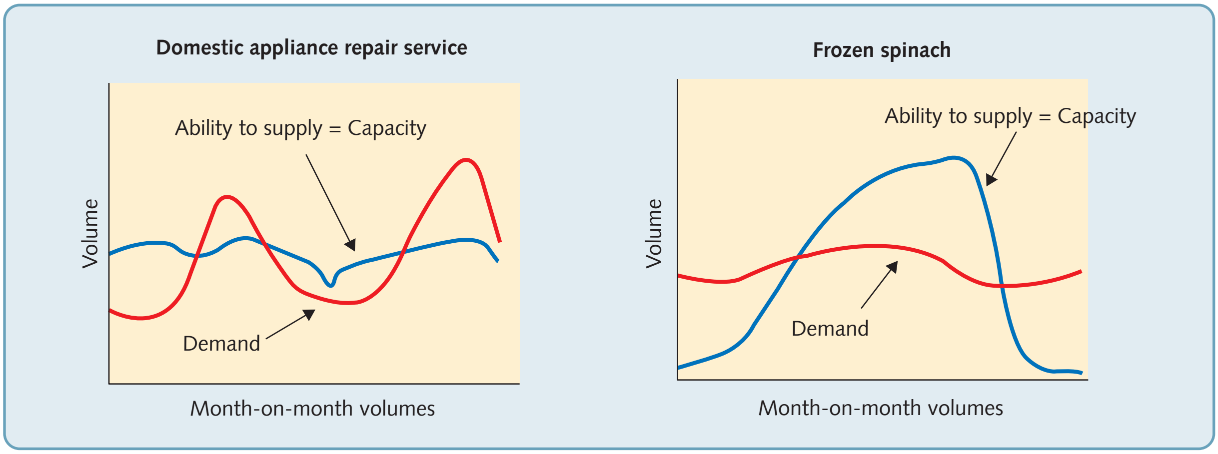

Figure 8.3 **Demand–capacity mismatches for an appliance repair service and a frozen spinach business**

> **OPERATIONS PRINCIPLE**
> ***Capacity management decision should reflect both predictable and unpredictable variations in capacity and demand.***

demand. The second business is a food manufacturer producing frozen spinach. The demand for this product is relatively constant throughout the year but the capacity of the business varies significantly. During the growing and harvesting season capacity to supply is high, but it falls off almost to zero for part of the year. Yet although the mismatch between demand and capacity is driven primarily by fluctuations in demand in the first case, and capacity in the second case, the essence of the capacity management activity is essentially similar for both.

DIAGNOSTIC QUESTION

What is the operation's current capacity?

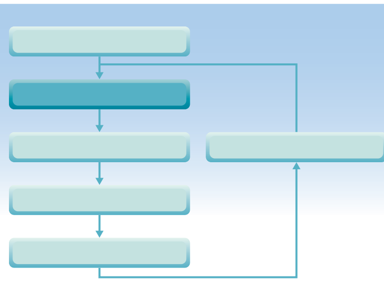

> **OPERATIONS PRINCIPLE**
> ***Measuring capacity is a function of product / service mix, duration and product service specification.***

The first step in managing capacity is to know what it actually is. Measuring capacity may sound simple, but can in fact be relatively hard to define unambiguously, unless the operation is standardised and repetitive. Any measure of capacity will contain a number of assumptions, each of which may be necessary in order to give an estimate of capacity, but each one obscures some aspect of reality. Again, taking capacity as 'the ability to supply', these assumptions relate to the mix of products or services supplied, the time over which they are supplied and the specification of what is supplied.

Capacity depends on product or service mix

How much an operation can do depends on what it is being required to do. For example, a hospital may have a problem in measuring its capacity because the nature of the products and service may vary significantly. If all its patients required relatively minor treatment with only short stays in hospital, it could treat many people per week. Alternatively, if most of its patients required long periods of observation or recuperation, it could treat far fewer. Output depends on the mix of activities in which the hospital is engaged and, because most hospitals perform many different types of activities, output is difficult (though not impossible!) to predict. Some of the problems caused by variation mix can be partially overcome by using aggregated capacity measures. 'Aggregated' means that different products and services are bundled together in order to get a broad view of demand and capacity. Medium-term capacity management is usually concerned with setting capacity levels in aggregated terms, rather than being concerned with the detail of individual products and services. Although this may mean some degree of approximation, especially if the mix of products or services being produced varies significantly, is usually acceptable, and is a widely used practice in medium-term capacity management. For example, a hotel might think of demand and capacity in terms of 'room nights per month'; this ignores the number of guests in each room and their individual requirements, but it is a good first approximation. A computer manufacturer might measure demand and capacity in the number of units it is capable of making per month, ignoring any variation in models.

Capacity depends on the duration over which output is required

Capacity is the output that an operation can deliver *in a defined unit of time.* The level of activity and output that may be achievable over short periods of time is not the same as the capacity that is sustainable on a regular basis. For example, a tax return processing office, during its peak periods at the end (or beginning) of the financial year, may be capable of processing 120,000 applications a week. It does this by extending the working hours of its staff, discouraging its staff from taking vacations during this period, avoiding any potential disruption to its IT systems (not allowing upgrades during this period, etc.), and maybe just by working hard and intensively. Nevertheless, staff do need vacations, nor can they work long hours continually, and eventually the information system will have to be upgraded. The capacity that is possible to cope with peak times is not sustainable over long periods. As such, capacity is taken to be the level of activity or output that can be sustained over an extended period time.

Capacity depends on the specification of output

Some operations can increase their output by changing the specification of the product or service (although this is more likely to apply to a service). For example, a postal service may effectively reduce its delivery dependability at peak times. So, during the busy Christmas period, the number of letters delivered the day after being posted may drop from 95 to 85 per cent. Similarly, accounting firms may avoid long 'relationship building' meetings with clients during busy periods. Important though these are, they can usually be deferred to less busy times. The important task is to distinguish between the 'must do' elements of the service that should not be sacrificed and the 'nice to do' parts of the service that can be omitted or delayed, in order to increase capacity in the short term.

EXAMPLE

Heathrow in a capacity crisis[3]

With over 469,000 flights and 72.3 million passengers arriving and departing each year, London Heathrow is one of the busiest international hubs in the world. And yet, it is an airport in crises. On an average day, 60 per cent of arrivals, totalling over 55,000 customers, spend time in one of Heathrow's four 'holding stacks'. These delays range from 4 to10 minutes, rising to 20 minutes in the late morning peak, when between 32 and 40 jets typically circle over London. The costs of these delays include £119,000 of wasted fuel per day, 600 tonnes of additional CO_2 emissions, and the frustration of many customers losing valuable work and leisure time. The key problem is operating capacity, which currently stands at 98 per cent – '[When] you have [one of] the most utilised pieces of infrastructure in the world, then one of the results is that you have airborne holding' said Jon Proudlove, managing director of the national air traffic service (NAS) at Heathrow. With no slack in capacity, the effect (as we have seen with the operations triangle earlier in this book) is that any variations (such as poor weather or poor weather on the ground) have an immediate impact on airplane processing speeds. The effects of Heathrow's capacity management problem are starting to be felt with several airlines seeking

alternative capacity in the UK and Europe to expand their operations. Yet, the solutions to the problem are far less clear. In the short and medium term, significant investments have been made to air traffic systems to try to increase existing capacity (i.e. to expand the 'ability to serve'), and improvements to boarding processes have been trialled to ensure rapid plane turn-around. However, in the longer term the capacity question inevitably turns to runway expansion. Heathrow currently operates two runways, compared with the four or five of its major European competitors: Schiphol, Madrid, Paris and Frankfurt. Yet, building a new runway, while perhaps the obvious *capacity* solution, has recently been vetoed by UK politicians in the face of strong, largely environmental, opposition. So for the time being, Heathrow's operations managers must manage existing capacity as best they can.

Capacity 'leakage'

Even after allowing for all the difficulties inherent in measuring capacity, the theoretical capacity of a process (the capacity that it was designed to have) is rarely achieved in practice. Some reasons for this are, to some extent, predictable. Different products or services may have different requirements, so people and machinery will have delays when switching between tasks. Maintenance will need to be performed on machines, while training will be required for employees. Scheduling difficulties could mean further lost time. Not all of these losses are necessarily avoidable; they may occur because of the market and technical demands on the process. However, some of the reduction in capacity can be the result of less predictable events. For example, labour shortages, quality problems, delays in the delivery of bought-in products and services, and machine, or system, breakdown, can all reduce capacity. This reduction in capacity is sometimes called 'capacity leakage' and one popular method of assessing this leakage is the overall equipment effectiveness (OEE) measure (see Figure 8.4) that is calculated as follows:

$$\text{OEE} = a \times p \times q$$

where a is the availability of a process, p is the performance or speed of a process and q is the quality of product or services that the process creates.

OEE works on the assumption that some capacity leakage occurs in the form of reduced availability. For example, availability can be lost through time losses such as set-up and changeover losses (when equipment, or people in a service context, are being prepared for the next activity), and breakdown failures (when the machine is being repaired or in a service context where employees are being trained / absent). Some capacity is lost through speed losses such as when equipment is idling (for example, when it is temporarily waiting for work from another process) and when equipment is being run below its optimum work rate. In a service context, the same principle can be seen when individuals are not working at an optimum rate; for example, mail order call centre employees in the quiet period after the winter holiday season. Finally, not everything processed by an operation will be error free. So some capacity is lost through quality losses.

For processes to operate effectively, they need to achieve high levels of performance against all three dimensions: availability, performance (speed) and quality. Viewed in isolation, these individual metrics are important indicators of performance, but they do not give a complete picture of the process's *overall* effectiveness. And critically, all these losses in the calculation mean that OEE represents the valuable operating time as a percentage of the capacity something was designed to have.

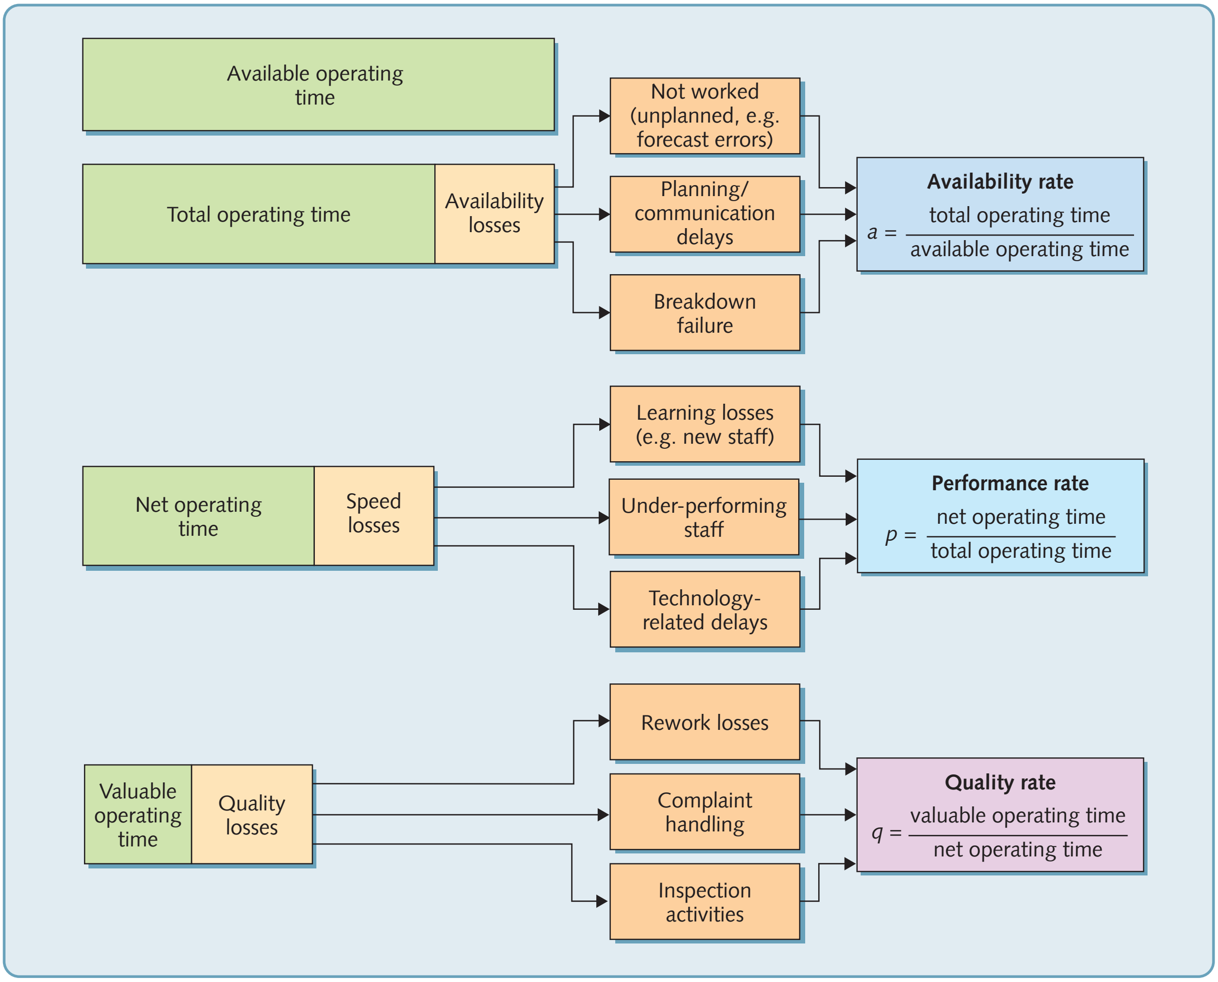

Figure 8.4 Overall equipment effectiveness

DIAGNOSTIC QUESTION

How well are demand–capacity mismatches understood?

Once demand and capacity are understood independently, managers must examine the potential mismatches that exist between them. For most organisations, this equates to understanding how demand might vary (although the same logic would apply to variation in capacity). In particular, the balance between predictable and unpredictable variation in demand affects the nature of capacity management. When demand is stable and predictable, the life of an operations manager is relatively easy! If demand is changeable, but this change is predictable, capacity adjustments may be needed, but at least they can be planned in advance. With unpredictable variation in demand, if an operation is to react to it at all, it must do so quickly; otherwise the change in capacity will have little effect on the operation's ability to deliver products and service as needed by their customers. Figure 8.5 illustrates how the objective and tasks of capacity management vary depending on the balance between predictable and unpredictable variation.

Enhanced market knowledge makes capacity planning easier

Capacity planning has to cope with mismatches between capacity and demand. Therefore, a deep understanding of the market forces that will generate demand is critical. This goes beyond the idea of forecasting as the prediction of uncontrollable events. Enhanced market knowledge is a broader concept and is also illustrated in Figure 8.5.

When the focus of capacity management is on dealing with predictable variation, demand–supply mismatches are, by definition, already known. What is useful under these circumstances is not so much knowledge of what demand–supply mismatches will be, but rather how they can be changed. So, for example, can a major customer be persuaded to move their demand to a quieter period? Will increasing prices at peak periods shift demand to off-peak periods? Can new storage techniques allow the supply of food ingredients throughout the year? However, when the demand–supply mismatch emerges largely through unpredictable events, then forecasting in its conventional sense is important because it converts unpredictable variation into predictable variation. Of course, forecasting cannot eliminate predictable variation, but it is a first step towards minimising the negative effects of variation by allowing better management of capacity (supply) and demand.

> OPERATIONS PRINCIPLE
> *Improving market knowledge allows an increased focus on predictable as opposed to unpredictable demand–capacity mismatches.*

Making forecasts useful for capacity management

Without some attempt to understand future demand and supply fluctuations, it is not possible to plan effectively for future events, only to react to them. But forecasts are (almost) always wrong (The economist John Kenneth Galbraith said, 'the only function of economic forecasting is to make astrology look respectable'). Yet it clearly helps capacity management if forecasts are

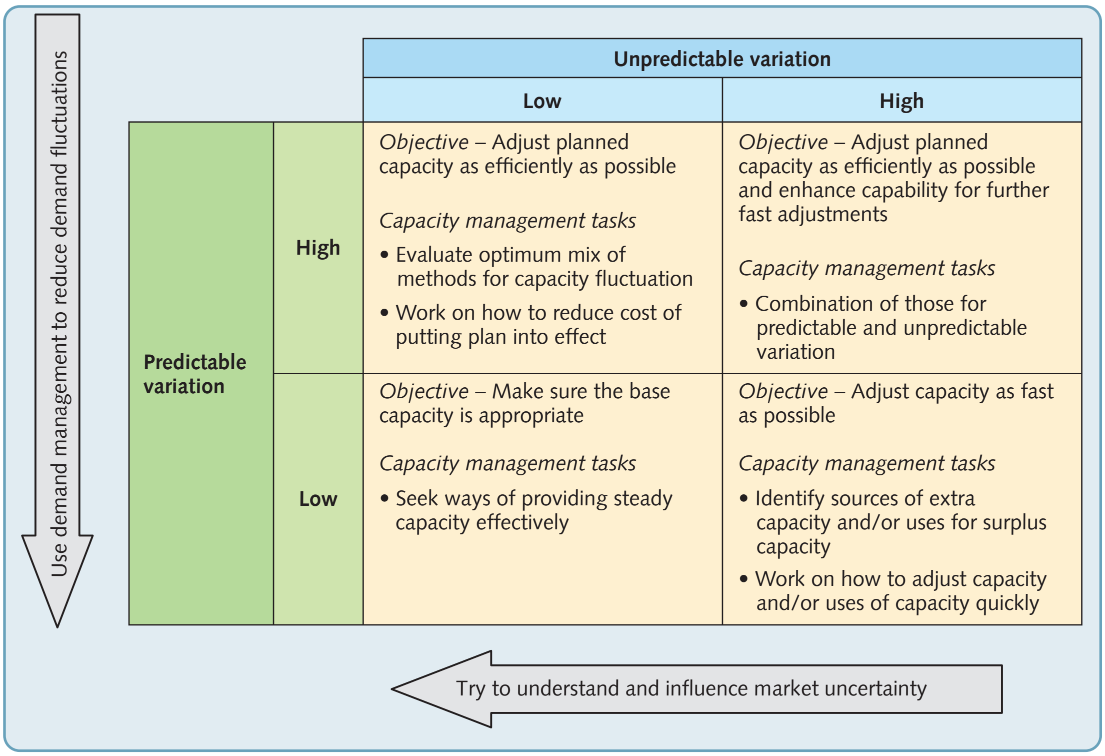

Figure 8.5 **The nature of capacity management depends on the mixture of predictable and unpredictable demand and capacity variation, but enhanced demand and supply market knowledge can make capacity management easier**

as accurate as possible (Forecasting was discussed in the supplement to Chapter 4). But, in addition to accuracy there are a number of other issues that make forecasts more (or less) useful as inputs to capacity planning.

Although forecasts can never be perfectly accurate all the time, sometimes, forecast errors are particularly damaging. For example, if a process is operating at a level close to its maximum capacity, over-optimistic forecasts could lead the process to committing itself to unnecessary capital expenditure to increase its capacity. Inaccurate forecasts for a process operating well below its capacity limit will also result in extra cost, but probably not to the same extent. So, critically, the effort put into forecasting should reflect the varying sensitivity to forecast error. Forecasts also need to be expressed in units that are useful for capacity planning. If forecasts are expressed only in money terms and give no indication of the demands that will be placed on an operation's capacity, they will need to be translated into realistic expectations of demand, expressed in the same units as the capacity. Nor should forecasts be expressed in money terms, such as sales, when those sales are themselves a consequence of capacity planning.

For example, some retail operations use sales forecasts to allocate staff hours throughout the day. Yet sales will also be a function of staff allocation. Better to use forecasts of 'traffic', the number of customers who potentially could want serving if there is sufficient staff to serve them. Perhaps most importantly, forecasts should give an indication of relative uncertainty. Demand in some periods is more uncertain than others. The importance of this is that the operations managers need an understanding of when increased uncertainty makes it necessary to have reserve capacity. A probabilistic forecast allows this type of judgement between possible plans that would virtually guarantee the operation's ability to meet actual demand, and plans that minimise costs. Ideally, this judgement should be influenced by the nature of the way the business wins orders: price-sensitive markets may require a risk-avoiding cost minimisation plan that does not always satisfy peak demand, whereas markets that value responsiveness and service quality may justify a more generous provision of operational capacity. Remember though, the idea that 'better forecasting' is needed for effective capacity management is only partly true. A better approach would be to say that it is enhanced market knowledge (both of demand and supply) generally that is important.

Better forecasting or better operations responsiveness?

The degree of effort (and cost) to devote to forecasting is often a source of heated debate within organisations. This often comes down to two opposing arguments. One goes something like this. *'Of course it is important for forecasts to be as accurate as possible; we cannot plan operations capacity otherwise. This invariably means we finish up with too much capacity (thereby increasing costs), or too little capacity (thereby losing revenue and dissatisfying customers).'* The counter argument is very different. *'Demand will always be uncertain, that is the nature of demand. Get used to it. The only way to satisfy customers is to make the operation sufficiently responsive to cope with demand, almost irrespective of what it is.'* Both these arguments have some merit, but both are extreme positions. In practice, operations must find some balance between having better forecasts and being able to cope without perfect forecasts.

OPERATIONS PRINCIPLE

Capacity management requires combining attempts to increase market knowledge with attempts to increase operations flexibility.

Trying to get forecasts right has particular value where the operation finds it difficult or impossible to react to unexpected demand fluctuations in the short term. Internet-based retailers at some holiday times, for example, find it difficult to flex the quantity of goods they have in stock in the short term. Customers may not be willing to wait. On the other hand, other types of operation working in intrinsically uncertain markets may develop fast and flexible processes to compensate for the difficulty in obtaining accurate forecasts. For example, fashion garment manufacturers try to overcome the uncertainty in their market by shortening their response time to new fashion ideas (cat walk to rack time) and the time taken to replenish stocks in the stores (replenishment time). Similarly, when the cost of not meeting demand is very high, processes also have to rely on their responsiveness rather than accurate forecasts. For example, accident and emergency departments in hospitals must be responsive, even if it means underutilised resources at times.

DIAGNOSTIC QUESTION

What should be the operation's base capacity?

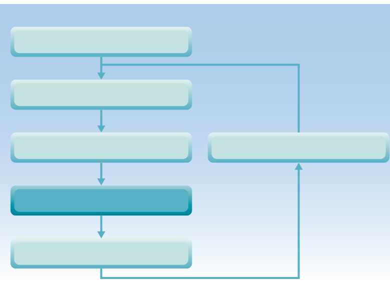

The most common way of planning capacity is to decide on a 'base level' of capacity and then adjust it periodically up or down to reflect fluctuations in demand. In fact, the concept of 'base' capacity is unusual because, although nominally it is the capacity level from which increases and decreases in capacity level are planned, in very unstable markets, where fluctuations are significant, it may never occur. Also, the two decisions of 'what should the base level of capacity be?' and 'how do we adjust capacity around that base to reflect demand?' are interrelated. An operation could set its base level of capacity at such a high level compared to demand that there is never a need to adjust capacity levels. However, this is clearly wasteful, which is why most operations will adjust their capacity level over time. Nevertheless, although the two decisions are interrelated, it is usually worthwhile setting a nominal base level of capacity before going on to consider how it can be adjusted.

OPERATIONS PRINCIPLE

The higher the base level of capacity, the less capacity fluctuation is needed to satisfy demand.

Setting base capacity

The base level of capacity in any operation is influenced by many factors, but should be related to three in particular:

- the relative importance of the operation's performance objectives
- the perishability of the operation's outputs
- the degree of variability in demand or supply.

Operation's performance objectives

Base levels of capacity should be set primarily to reflect an operation's performance objectives (see Figure 8.6). For example, setting the base level of capacity high compared to average demand will result in relatively high levels of underutilisation of capacity and therefore high costs. This is especially true when an operation's fixed costs are high and therefore the consequences of underutilisation are also high. Conversely, high base levels of capacity result in a capacity 'cushion' for much of the time, so the ability to flex output to give responsive customer service will be enhanced. When the output from the operation is capable of being stored, there may also be a trade-off between fixed capital and working capital in where base capacity level is set. A high level of base capacity can require considerable investment, while a lower base level would reduce the need for capital investment but may require inventory to be built up to satisfy future demand, thus increasing working capital. For some operations, building up inventory is either risky because products have a short

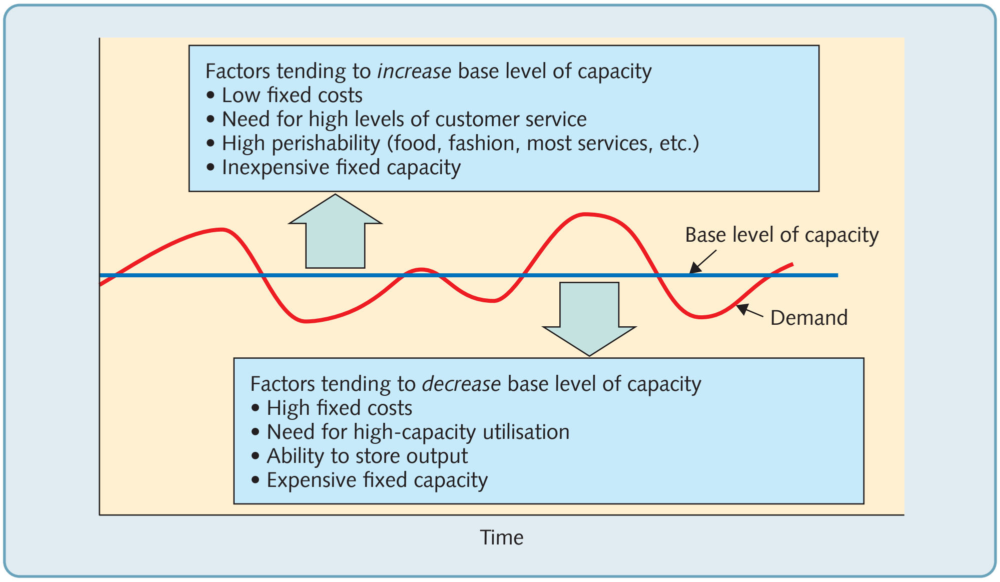

Figure 8.6 The base level of capacity should reflect the relative importance of the operation's performance objectives

shelf-life (for example, perishable food, high-performance computers, or fashion items) or because the output cannot be stored at all (most services).

The perishability of the operation's outputs

When either supply or demand is perishable, base capacity will need to be set at a relatively high level because inputs to the operation or outputs from the operation cannot be stored for long periods. For example, a factory that produces frozen fruit will need sufficient freezing, packing and storage capacity to cope with the rate at which the fruit crop is harvested during its harvesting season. Similarly, a hotel cannot store its accommodation services. If an individual hotel room remains unoccupied, the ability to sell for that night has 'perished'. In fact, unless a hotel is fully occupied every single night, its capacity is always going to be higher than the average demand for its services.

The degree of variability in demand or supply

Variability, either in demand or capacity will reduce the ability of an operation to process its inputs. That is, it will reduce its effective capacity. This effect was explained in Chapter 6 when the consequences of variability in individual processes were discussed. As a reminder, the greater the variability in arrival time or activity time at a process, the more the process will suffer both high throughput times *and* reduced utilisation. This principle holds true for whole operations, and because long throughout put times mean that queues will build up in the operation, high variability also affects inventory levels. This is illustrated in Figure 8.7. The implication of this is that the greater the variability, the more extra capacity will need to be provided to compensate for the reduced utilisation of available capacity. Therefore, operations with high levels of variability will tend to set their base level of capacity relatively high, in order to provide this extra capacity. Of course, as we have seen earlier, not all operations have the option of simply increasing capacity! (See the example, 'Heathrow in capacity crisis'.)

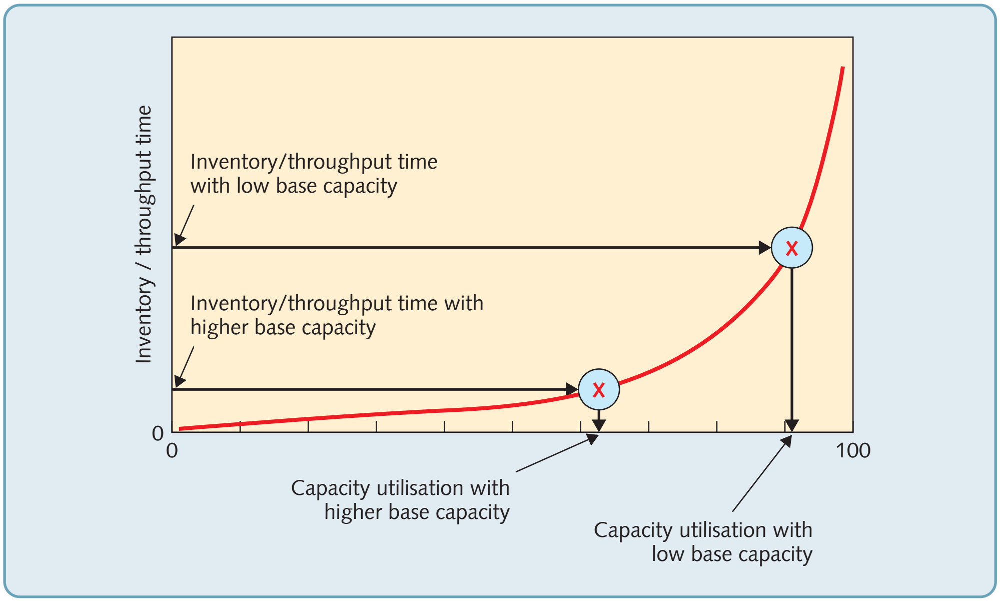

Figure 8.7 **The effect of variability on the utilisation of capacity**

DIAGNOSTIC QUESTION

How can demand–capacity mismatches be managed?

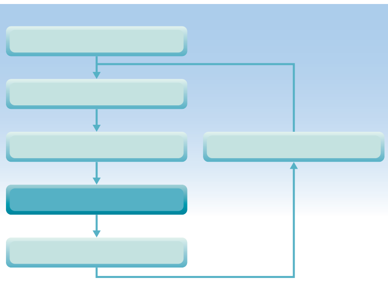

Almost all operations have to cope with varying demand or supply; therefore, they will need to consider adjusting capacity around its nominal base level. There are three 'pure' plans available for treating such variation, although in practice, most organisations will use a mixture of all of them, even if one plan dominates:

1. *Level capacity plan:* ignore demand fluctuations and keep nominal capacity levels constant.
2. *Chase demand plan:* adjust capacity to reflect the fluctuations in demand.
3. *Demand management:* attempt to change demand.

Level capacity plan

In a level capacity plan, the capacity is fixed throughout the planning period, regardless of the fluctuations in forecast demand. This means that the same number of staff or machines operate the same processes and should therefore be capable of producing the same aggregate output in each period. Where non-perishable materials are processed, but not immediately sold, they can be transferred to finished goods inventory in anticipation of later sales. When inventory is not possible, as in most service operations, demand fluctuations are absorbed through underutilisation of the operation's resources and/or and failure to meet demand immediately (see Figure 8.8(a)). The more demand fluctuates, the higher the level of inventory or underutilisation there is. Both are expensive, but may be considered if the cost of building inventory is low compared with

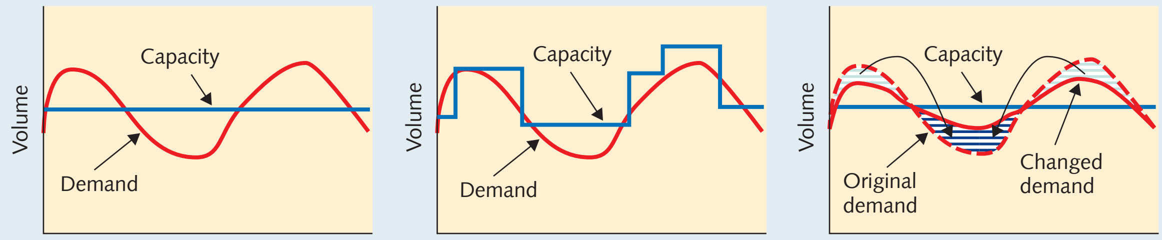

(a) Level capacity plan – absorb fluctuations

(b) Chase demand plan – change capacity to reflect demand fluctuations

(c) Manage demand plan – attempt to change demand to reduce fluctuations

Figure 8.8 Managing demand–capacity mismatches using 'level capacity', 'chase demand' and 'demand management' plans

changing output levels, or in service operations, if the opportunity costs of individual lost sales is very high: for example, in the high-margin retailing of jewellery and in (real) estate agents. Setting capacity below the forecast peak demand level will reduce the degree of underutilisation, but, in the periods where demand is expected to exceed capacity, customer service may deteriorate

Chase demand plan

Chase demand plans attempt to match capacity closely to the varying levels of forecast demand, as in Figure 8.8(b). This is much more difficult to achieve than a level capacity plan, so for this reason, pure chase demand plans are unlikely to appeal to operations producing standard, non-perishable products, especially where operations are capital-intensive. The chase demand policy would require a level of physical capacity (as opposed to effective capacity), all of which would only be used occasionally. A pure chase demand plan is more usually adopted by operations that cannot store their output, such as a call centre. It avoids the wasteful provision of excess staff that occurs with a level capacity plan, and yet should satisfy customer demand throughout the planned period. Where inventory is expensive to hold, a chase demand policy might be adopted in order to minimise it. There are a number of different methods for adjusting capacity, although they may not all be feasible for all types of operation. Some of these methods are shown in Table 8.1.

EXAMPLE

Annualised hours at Lowaters[4]

Lowaters Nursery is a garden plant and horticulture specialist in the South of England, employing around 25 people. Like any business that depends on seasonal weather conditions, it faces fluctuating demand for its services and products. It also prides itself on offering *'the best service in partnership with our customers, by communicating in a friendly professional manner and listening to our customers to provide the result required'* (Lowaters' mission statement). But to maintain its quality of service throughout the seasonal ups and down in workload means keeping your core team happy and employed throughout the year. This is why Lowaters introduced its annualised hours scheme, a method of

Table 8.1 Summary of advantages and disadvantages of some methods of adjusting capacity

Method of adjusting capacity	*Advantages*	*Disadvantages*
Overtime – staff working longer than their normal working times	Quickest and most convenient	Extra payment normally necessary and agreement of staff to work, can reduce productivity over long periods
Annualised hours – staff contracting to work a set number of hours per year rather than a set number of hours per week	Without many of the costs associated with overtime the number of staff time available to an organisation can be varied throughout the year to reflect demand	When very large and unexpected fluctuations in demand are possible, all the negotiated annual working time flexibility can be used before the end of the year
Staff scheduling – arranging working times (start and finish times) to vary the aggregate number of staff available for working at any time	Staffing levels can be adjusted to meet demand without changing job responsibilities or hiring in new staff	Providing start and finish (shift) times that both satisfy staffs' need for reasonable working times and shift patterns as well as providing appropriate capacity can be difficult
Varying the size of the workforce – hiring extra staff during periods of high demand and laying them off as demand falls, or hire and fire	Reduces basic labour costs quickly	Hiring costs and possible low productivity while new staff go through the learning curve. Lay-offs may result in severance payments and possible loss of morale in the operation and loss of goodwill in the local labour market
Using part-time staff – recruit staff who work for less than the normal working day (at the busiest periods)	Good method of adjusting capacity to meet predictable short-term demand fluctuations	Expensive if the fixed costs of employment for each employee (irrespective of how long he or she works) are high
Skills flexibility – designing flexibility in job design and job demarcation so that staff can transfer across from less busy parts of the operation	Fast method of reacting to short-term demand fluctuations	Investment in skills training needed and may cause some internal disruption
Subcontracting / outsourcing – buying, renting or sharing capacity or output from other operations	No disruption to the operation	Can be very expensive because of subcontractor's margin and subcontractor may not be as motivated to give same service, or quality. Also a risk of leakage of knowledge
Change output rate – expecting staff (and equipment) to work faster than normal	No need to provide extra resources	Can only be used as a temporary measure, and even then can cause staff dissatisfaction, a reduction in the quality of work, or both

fluctuating capacity as demand varies throughout the year without many of the costs associated with overtime or hiring temporary staff. It involves staff contracting to work a set number of hours per year, rather than a set number of hours per week. The main advantage of this is that the amount of staff time available to an organisation can be varied throughout the year to reflect the real state of demand. Annual hours plans can also be useful when supply varies throughout the year. Maria Fox, one of the management team at Lowaters says that annualised hours give the company several advantages. '*It simplifies administration and gives us the flexibility we need to run the business, while delivering some real advantages to the employees. They are all effectively on salary with fixed monthly payments. We can flex the hours worked over the year – when we are busy we work longer and when things are quiet, in the winter, they can take time off. Everyone other than directors is contracted to work 39 hours on average over 52 weeks of the year.*

The company created a simple spreadsheet that sets out the actual hours worked and compares them with a target distribution of the annualised hours that are expected to be worked over the year. This allows employees to see at a glance whether someone is over or under target. '*We email them a copy of their sheet at the beginning of the year so they can keep track of their own progress as they go*', said Maria Fox. '*It also allows us to keep track of how many hours they do. If at the end of the year they come in plus or minus 50 hours, we simply adjust it up or down for the next year. If there is a bigger discrepancy than that, we'll look at the job structure.*'

However, not all experiments with annualised hours have been as successful as that at Lowaters. Where demand is very unpredictable, staff can be asked to come in to work at very short notice. This can cause considerable disruption to social and family life. For example, at one news-broadcasting company, the scheme caused problems. Journalists and camera crew who went to cover a foreign crisis found that they had worked so many hours, they were asked to take the whole of one month off to compensate. Since they had no holiday plans, many would have preferred to work for additional income

Zero-hours contracts

A zero-hours contract is one where an employer does not offer any guarantee of a specific number of hours of work for a worker. Neither is any person working under this type of contract under any obligation to accept those hours when they are offered. It can therefore be seen as a method of varying effective capacity for 'chase demand' approaches. The prevalence of zero-hours arrangements, exactly how they are defined and their legal status varies significantly between countries. For example, at the time of writing they are allowed in Sweden, Norway and the UK, allowed but regulated in Germany and the Netherlands, and not generally allowed in France and Austria.[5] However, they are controversial, even where they are not particularly widely used. Under such schemes, people may have no idea about how many hours they will be working (and therefore the wages that they will be earning). But the flexibility may be valued by some people, notwithstanding the fluctuating incomes. Certainly some employers cite the flexibility that zero-hours contracts gives their operations as a major advantage.

Changing capacity when variation is unpredictable

Both the mix of methods used to change capacity and how they are implemented will depend on the balance between predictable and unpredictable variation. As we discussed earlier, the objective of capacity management when demand variation is predictable is to affect the changes as efficiently as possible. Whereas when demand fluctuations are unpredictable, the objective is usually to change capacity as fast as possible. In the latter case, it is necessary to understand the flexibility of the resources that may be used to increase capacity. In this case we are using flexibility to mean both how much capacity can be changed and how fast it can be changed. In fact, the degree of change and the response time required to make the change are almost always related. The relationship can be shown in what is termed a 'range–response' curve. Figure 8.9 shows one of these for a call centre. It shows that within a few minutes of demand for the call centre's services increasing, it has the ability to switch a proportion of its calls to the company's other call centres. However, not everyone in these other call centres is trained to take such calls, therefore any further increase in capacity must come from bringing in staff currently not on shift. Eventually, the call centre will hit its limits of physical capacity (computers, telephone lines, etc.). Any further capacity increase will have to wait until more physical capacity is added.

Demand management

The objective of demand management is to change the pattern of demand to bring it closer to available capacity, usually by transferring customer demand from peak periods to quiet periods, as was shown in Figure 8.8(c). There are a number of methods for achieving this:

- **Constraining customer access** – customers may only be allowed access to the operation's products or services at particular times. For example, reservation and appointment systems in hospitals.
- **Price differentials** – adjusting price to reflect demand; that is, increasing prices during periods of high demand and reducing prices during periods of low demand.

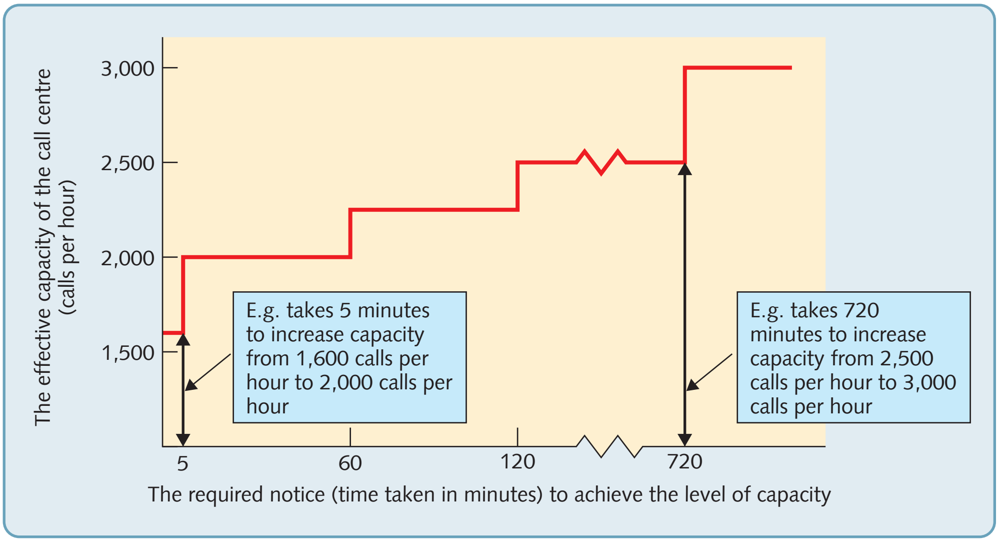

Figure 8.9 **The range–response curve for increasing capacity at a call centre**

- **Scheduling promotion** – varying the degree of market stimulation through promotion and advertising, in order to encourage demand during normally low periods.
- **Service differentials** – allowing service levels to reflect demand (implicitly or explicitly), allowing service to deteriorate in periods of high demand and increase in periods of low demand. If this strategy is used explicitly, customers are being educated to expect varying levels of service and hopefully move to periods of lower demand.

A more radical approach attempts to create alternative products or services to fill capacity in quiet periods. It can be an effective demand management method but, ideally, new products or services should meet three criteria: (a) they can be produced on the same processes; (b) they have different demand patterns to existing offerings; and (c) they are sold through similar marketing channels. For example, ski resorts may provide organised mountain activity holidays in the summer, and garden tractor companies may make snow movers in the autumn and winter. However, the apparent benefits of filling capacity in this way must be weighed against the risks of damaging the core product or service, and the operation must be fully capable of serving both markets.

EXAMPLE

The surge-pricing dilemma[6]

Surge (or dynamic) pricing is a demand management technique that relies on frequent adjustments in price to influence supply and (especially) demand so that they match each other. For example, some electricity suppliers charge different rates for energy depending on when it is consumed. Similarly, in countries with road charging, tolls are set at higher levels during peak times in an effort to keep traffic flowing. But perhaps the best-known example of surge pricing is the algorithm used by the taxi app Uber. During times of excessive demand or inadequate supply, when the number of people wanting a ride exceeds the number of available drivers, Uber applies a multiplier to increase its normal fares based on the scarcity of available drivers. Uber said that it does this, '*to make sure those who need a ride can get one. . . [and] . . . for riders, surge [pricing] helps ensure that pickup is available quickly and reliably. . . for driver-partners, surge means higher fares and a steady stream of ride requests*'. The problem with surge pricing is that it is efficient yet deeply unpopular with customers. Economists may understand and have faith in the power of supply and demand, but most of Uber's customers are not economists. In the press and on social media, customers complain that it seems that the company are taking advantage of them. But some marketing experts say it is, at least partly, a matter of perception. As well as perhaps capping their multiplier, Uber should make the way they calculate it more transparent, limit how often prices are adjusted, communicate the benefits of the technique, and change its name (certainty pricing and priority pricing have been suggested).

Yield management

In operations that have relatively fixed capacities, such as airlines and hotels, it is important to use the capacity of the operation for generating revenue to its full potential. One approach used by such operations is called yield management.[7] This is really a collection of methods, some of which we have already discussed, which can be used to ensure that an operation maximises its potential to generate profit. Yield management is especially useful where capacity is relatively fixed; the market can be fairly clearly segmented; the service cannot be stored in any way; the service is sold in advance; and the marginal cost of making a sale is relatively low.

Airlines, for example, fit all these criteria. They adopt a collection of methods to try to maximise the yield (i.e. profit) from their capacity. Over-booking capacity may be used to compensate for passengers who do not show up for the flight. However, if more passengers show up than they expect, the airline will have a number of upset passengers (although they may be able to offer financial inducements for the passengers to take another flight). By studying past data on flight demand, airlines try to balance the risks of over-booking and under-booking. Operations may also use price discounting at quiet times, when demand is unlikely to fill capacity. For example, hotels will typically offer cheaper room rates outside of holiday periods to try and increase naturally lower demand. In addition, many larger chains will sell heavily discounted rooms to third parties who in turn take on the risk (and reward) of finding customers for these rooms.

EXAMPLE

United drags passengers off its plane[8]

When footage shot by a fellow passenger showed a bloodied and unconscious man being pulled off of a flight at O'Hare international airport in Chicago, the clip caused a sensation on social media and later mainstream news outlets. It also resulted in a barrage of criticism on United Airlines, the offending carrier. The incident began when United overbooked the flight (a problem made worse because at the last minute it decided to fly four members of staff to a connection point so they could staff another flight). The airline decided that it needed to bump four passengers to make way for them. This is common practice; when a flight is overbooked, the first step is to offer a financial encouragement for passengers to take a later flight. In this case, passengers were first offered $400, overnight hotel accommodation and a flight the following day. No one accepted and the offer was increased to $800. Still no one accepted the offer, so a manager announced that passengers would be selected to leave the flight, with frequent fliers and business-class passengers being given priority. The first two people selected agreed to leave the plane. The third person selected (as it happened, the wife of the man who was later dragged off forcibly), also agreed. However, when the fourth person was approached, the man refused, saying that he was a doctor and had to see patients in the morning. Eyewitnesses said the man was 'very upset' and tried to call his lawyer. So, instead of selecting another passenger, or increasing its offer (it could have offered a maximum of $1,350), security staff were called. The encounter with the security staff concluded with the man being wrenched from his seat onto the floor, after which he was hauled down the aisle, blood covering his face.

It is not uncommon for airlines to sell more tickets than they have seats, under the assumption that some passengers will either fail to show, or cancel at the last minute. It is claimed that because of fierce competition and to avoid flying with empty seats, so-called 'denied boarding' incidents are becoming more common. As the incident gained publicity, the CEO of United Airlines, Oscar Munoz, said that employees '*followed established procedures', but he was 'upset to see and hear about what happened*', although the passenger had been '*disruptive and belligerent*'. When the passenger refused to voluntarily leave the plane, he said, the staff were '*left with no choice but to call security officers to assist in removing the customer from the flight*'. Travel expert, Simon Calder, said that the airline was technically within their rights. '*The captain is in charge of the aircraft. And if he or she decides that someone needs to be offloaded, that command has to be obeyed. From the moment that the unfortunate individual in this case said, "I'm staying put", he became a disruptive passenger. Officials were legally entitled to remove him, and as the videos show, he was dragged from the plane. It appears from the evidence that the law was broken – by him, not by the airline. But I would be surprised if United pressed charges.*'

Using cumulative representations to plan capacity

When an operation's output can be stored (in contrast to yield management, where it cannot), a useful method of assessing the feasibility and consequences of adopting alternative capacity plans is the use of cumulative demand and supply curves. These plot (or calculate) both the cumulative demand on an operation, and its cumulative ability to supply, over time. For example, Figure 8.10 shows the forecast aggregated demand for a chocolate factory which makes confectionery products. Demand for its products in the shops is greatest in December. To meet this demand and allow time for the products to work their way through the supply chain, the factory must supply a demand which peaks in September. But the cumulative representation of demand against available supply time (productive days) shown in Figure 8.10 reveals that although total demand peaks in September, because of the restricted number of available productive days, the peak demand per productive day occurs a month earlier in August. It also shows that the effective fluctuation in demand over the year is even greater than it seemed. The ratio of monthly peak demand to monthly lowest demand is 6.5:1, but the ratio of peak to lowest demand per productive day is 10:1. Demand per productive day is more relevant to operations managers, because productive days represent the 'ability to supply'.

The feasibility and consequences of a capacity plan can be assessed on this basis. Figure 8.10 also shows a level capacity plan (A) that assumes production at a rate of 14.03 tonnes per productive day. This meets cumulative demand by the end of the year, so there is theoretically no over-capacity or under-capacity. However, if one of the aims of the plan is to supply demand when it occurs, the plan is inadequate. Up to around day 168, the line representing cumulative production is above that representing cumulative demand. This means that at any time during this period, more product has been produced by the factory than has been demanded from it. In fact the vertical distance between the two lines is the level of inventory at that point in time. So by day 80, 1,122 tonnes have been produced but only 575 tonnes have been demanded. The surplus of production above demand, or inventory, is therefore 547 tonnes. When the cumulative demand line lies above the cumulative production line, the reverse is true. The vertical distance between the two lines now indicates the shortage, or lack of supply. So by day 198, 3,025 tonnes have been demanded but only 2,778 tonnes produced. The shortage is therefore 247 tonnes.

For any capacity plan to meet demand as it occurs, its cumulative production line must always lie above the cumulative demand line. This makes it a straightforward task to judge the adequacy of a plan, simply by looking at its cumulative representation. An impression of the inventory implications can also be gained from a cumulative representation by judging the area between the cumulative production and demand curves. This represents the amount of inventory carried over the period. Level capacity plan (B) is feasible because it always ensures enough production to meet demand at any time throughout the year. However, inventory levels are high using this plan. It may even mean that the chocolate spends so much time in the factory's inventory, that it has insufficient shelf-life when it arrives at the company's retail customers. Assuming a 'first-in-first-out' inventory management

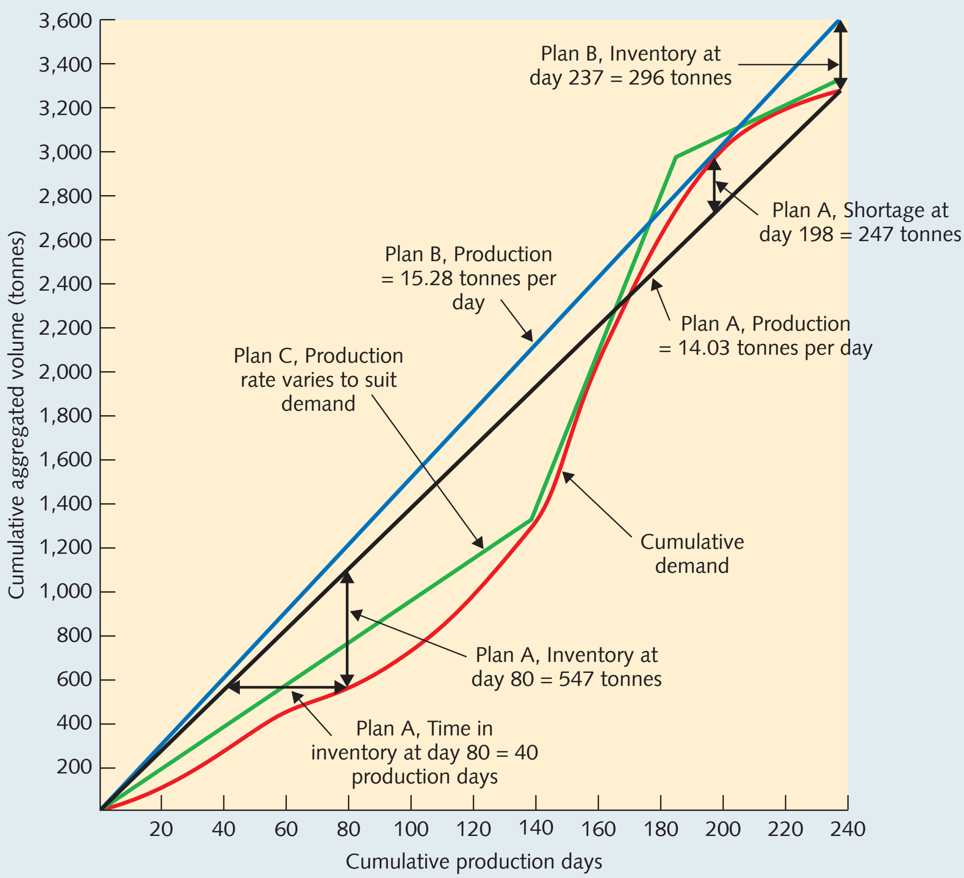

	J	F	M	A	M	J	J	A	S	O	N	D
Demand (tonnes/month)	100	150	175	150	200	300	350	500	650	450	200	100
Production days	20	18	21	21	22	22	21	10	21	22	21	18
Demand (tonnes/day)	5	8.33	8.33	7.14	9.52	13.64	16.67	50	30.95	20.46	9.52	5.56
Cumulative days	20	38	59	80	102	124	145	155	176	198	219	237
Cumulative demand	100	250	425	575	775	1075	1425	1925	2575	3025	3225	3325

Figure 8.10 Cumulative representation of demand and three capacity plans

OPERATIONS PRINCIPLE

For any capacity plan to meet demand as it occurs, its cumulative production line must always lie above its cumulative demand line.

principle, the time the product stays in inventory will be represented by the horizontal line between the demand at the time it is 'demanded' and the time it was produced.

Inventory levels (and therefore the time products spend as part of the inventory) can be reduced by adopting a chase demand plan, such as that shown as (C) in Figure 8.10. This reduces inventory carrying costs, but incurs costs associated with changing capacity levels. Usually, the marginal cost of making a capacity change increases with the size of the change. For example, if the chocolate manufacturer wishes to increase capacity by 5 per cent, this can be achieved by requesting its staff or work overtime – a simple, fast and relatively inexpensive option. If the change is 15 per cent, overtime cannot provide sufficient extra capacity and temporary staff will need to be employed – a more expensive solution which also would take more time. Increases in capacity of above 15 per cent might only be achieved by subcontracting some work out. This would be even more expensive.

DIAGNOSTIC QUESTION

How should capacity be controlled?

Although planning capacity levels in advance, and even planning how to respond to unexpected changes in demand, is an important part of capacity management, it does not fully reflect the dynamic nature of the activity. Capacity management must react to *actual* demand and *actual* capacity as it occurs. Period by period, operations management considers its forecasts of demand, its understanding of current capacity and, if outputs can be stocked, how much inventory (or in service contexts, queues of customers in the process) has been carried forward from the previous period. Based on all this information, it makes plans for the following period's capacity. During the next period, demand might or might not be as forecast and the actual capacity of the operation might or might not turn out as planned (because of the capacity leakage discussed earlier). But whatever the actual conditions during that period, at the beginning of the next period the same types of decisions must be made, in the light of the new circumstances. Figure 8.11 shows how this works in practice. It shows the overall performance of an operation's capacity management as a function of the way it manages capacity and the way it manages (or forecasts) demand.

The success of capacity management is generally measured by some combination of costs, revenue, working capital and customer satisfaction (which goes on to influence revenue). This is influenced by the actual capacity available to the operation in any period and the demand for that period. If capacity is in excess of demand, customer demands can be met, but

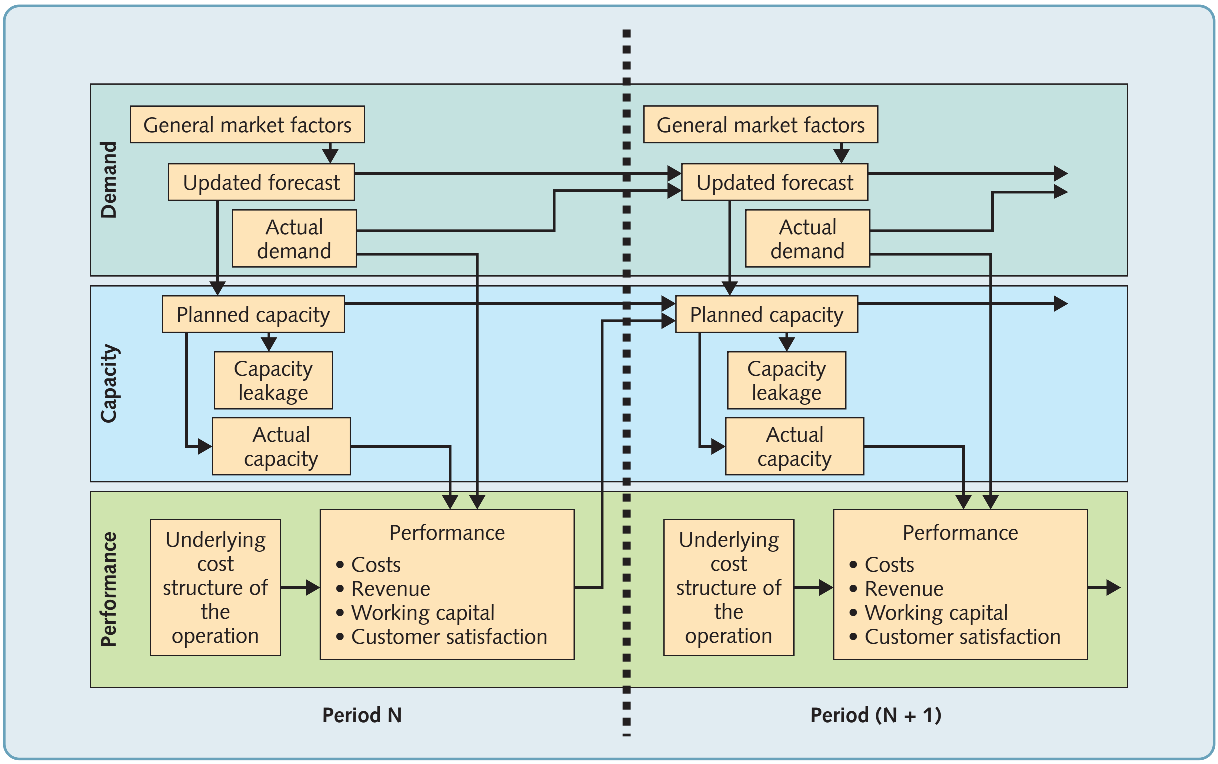

Figure 8.11 How should capacity be controlled – the dynamics of capacity management

Short-term outlook for volume

Long-term outlook for volume	Decreasing below current capacity	Level with current capacity	Increasing above current capacity
Decreasing below current capacity	Reduce capacity (semi) permanently. For example, reduce staffing levels; reduce supply agreements.	Plan to reduce capacity (semi) permanently.For example, freeze recruitment; modify supply agreements.	Increase capacity temporarily. For example, increase working hours, and/or hire temporary staff; modify supply agreements.
Level with current capacity	Reduce capacity temporarily. For example, reduce staff working hours; modify supply agreements.	Maintain capacity at current level.	Increase capacity temporarily. For example, increase working hours, and/or hire temporary staff; modify supply agreements.
Increasing above current capacity	Reduce capacity temporarily. For example, reduce staff working hours, but plan to recruit; modify supply agreements.	Plan to increase capacity above current level; plan to increase supply agreements.	Increase capacity (semi) permanently. For example, hire staff; increase supply agreements.

Figure 8.12 Capacity management strategies are partly dependent on the long and short-term outlook for volumes

underutilised capacity and possibly inventory, will increase costs. If capacity is less than demand, the operation's resources will be fully utilised, but at the expense of being unable to meet all demand. The effect is that customers will be waiting for products and service (or in many cases will be switching to a competitor who is able to provide what they need). But, some operations are more able to cope than others with any mismatch between actual capacity and actual demand. If the underlying cost structure of the operation is such that fluctuations in output level have relatively little effect on costs, then the operation will be less sensitive to errors in capacity management. However, overriding other considerations of what capacity strategy to adopt is often any difference between the long-and short-term outlook for the volume of demand. See Figure 8.12.

Demand forecasting should always be an ongoing process that incorporates the general market factors that influence demand. In addition, the actual demand that occurs each month should be factored into each period's forecast. In doing so, the operation should aim to build up experience of managing demand, managing capacity, and adapting the operation to make it less sensitive to mismatches between the two. Successful capacity control also requires businesses to learn from their handling of previous demand fluctuations. Period by period, operations managers are reacting to a set of stimuli as illustrated in Figure 8.12. Some of these stimuli may be ambiguous, such as the overall objectives of the operation and its approach to risk. Others will be uncertain, such as future demand and to a lesser extent future capacity. This is a complex decision-making process that depends on more than the availability and accuracy of information. It also depends on the ability to refine decision-making behaviour through learning from past successes and mistakes. For example, some managers may have the tendency to overreact to immediate stimuli by frequently increasing or decreasing

OPERATIONS PRINCIPLE

The learning from managing capacity in practice should be captured and used to refine both demand forecasting and capacity planning.

capacity as forecasts of future demand are adjusted. These overreactions can relate to personality traits, previous experience, different attitudes to downside and upside risk, or industry norms. Whatever the reason, some mechanism is needed for smoothing both forecasts and the response to them.

Critical commentary

For such an important topic, there is surprisingly little standardisation in how capacity is measured. Not only is a reasonably accurate measure of capacity needed for operations management, it is also needed to decide whether it is worth investing in extra physical capacity. Yet not all practitioners would agree with the way in which capacity has been defined or measured in this chapter (although it does represent orthodox practice). One school of thought is that whatever capacity efficiency measures are used, they should be useful as diagnostic measures which can highlight the root causes of inefficient use of capacity. The idea of overall equipment effectiveness (OEE) described earlier is often put forward as a useful way of measuring capacity efficiencies.

- The other main point of controversy in capacity management concerns the use of varying staff levels. To many, the idea of fluctuating the workforce to match demand, by using part-time staff, or using zero-hours contracts, or by hiring and firing, is more than just controversial. It is regarded as unethical. It is any business's responsibility, they argue, to engage in a set of activities that are capable of sustaining employment at a steady level. Hiring and firing merely for seasonal fluctuations, which can be predicted in advance, is treating human beings in an unacceptable manner. Even hiring people on a short-term contract, in practice, leads to them being offered poorer conditions of service and leads to a state of permanent anxiety as to whether they will keep their jobs. On a more practical note, it is pointed out that in an increasingly global business world, where companies may have sites in different countries, those countries that allow hiring and firing are more likely to have their plants 'downsized' than those where legislation makes this difficult.

SUMMARY CHECKLIST

- ☐ Is the importance of effective capacity management fully understood?
- ☐ Is the operation's current capacity measured?
- ☐ If so, are all the assumptions inherent in the measurement of capacity made fully explicit?
- ☐ What capacity 'leakage' is normal, and have options for minimising capacity leakage been explored?
- ☐ Is there scope for using the overall equipment effectiveness (OEE) measure of capacity?
- ☐ What is the balance between predictable variation and unpredictable variation in demand and capacity?
- ☐ Realistically, what potential is there for making unpredictable variability more predictable through better forecasting?
- ☐ Does an understanding of the market include the extent to which the behaviour of customers and/or suppliers can be influenced to reduce variability?
- ☐ Does the operations base capacity reflect all the factors that should be influencing its level?
- ☐ Have alternative methods of adjusting (or not) capacity been fully explored and assessed?
- ☐ If variation is unpredictable, have methods of speeding up the operation's reaction to demand–capacity mismatches been explored?
- ☐ Is there scope for using cumulative representations of demand and capacity for planning purposes?
- ☐ Is the method of deciding period-by-period capacity levels effective?
- ☐ How does the method of deciding period-by-period capacity levels reflect previous experience?

CASE STUDY

Blackberry Hill Farm

'Six years ago I had never heard of agri-tourism. As far as I was concerned, I had inherited the farm and I would be a farmer all my life.'(Jim Walker, Blackberry Hill Farm)

The 'agri-tourism' that Jim was referring to is 'a commercial enterprise at a working farm, or other agricultural centre, conducted for the enjoyment of visitors that generates supplemental income for the owner'. *'Farming has become a tough business,'* said Jim. *'Low world prices, a reduction in subsidies, and increasingly uncertain weather patterns have made it a far more risky business than when I first inherited the farm. Yet, because of our move into the tourist trade we are flourishing. Also. . . I've never had so much fun in my life'.* But, Jim warns, agri-tourism isn't for everyone. *'You have to think carefully. Do you really want to do it? What kind of life style do you want? How open-minded are you to new ideas? How business-minded are you? Are you willing to put a lot of effort into marketing your business? Above all, do you like working with people? If you had rather be around cows than people, it isn't the business for you.'*

History

Blackberry Hill Farm was a 200 hectare mixed farm in the South of England when Jim and Mandy Walker inherited it 15 years ago. It was primarily a cereal growing operation, with a small dairy herd, some fruit and vegetable growing and mixed woodland that was protected by local preservation laws. Six years ago it had become evident to Jim and Mandy that they may have to rethink how the farm was being managed. *'We first started a pick-your-own (PYO) operation because our farm is close to several large centres of population. Also the quantities of fruit and vegetables that we were producing were not large enough to interest the commercial buyers. Entering the PYO market was a reasonable success and in spite of making some early mistakes, it turned our fruit and vegetable growing operation from making a small loss to making a small profit. Most importantly, it gave us some experience of how to deal with customers face to face and of how to cope with unpredictable demand. The biggest variable in PYO sales is weather. Most business occurs at the weekends between late spring and early autumn. If rain keeps customers away during part of those weekends, nearly all sales have to occur in just a few days.'*

Within a year of opening up the PYO operation Jim and Mandy had decided to reduce the area devoted to cereals and increase their fruit and vegetable growing capability. At the same time they organised a 'Petting Zoo' that allowed children to mix with, feed and touch various animals.

'We already had our own cattle and poultry but we extended the area and brought in pigs and goats. Later we also introduced some rabbits, ponies and donkeys, and even a small bee-keeping operation.' At the same time, the farm started building up its collection of 'farm heritage' exhibits. These were static displays of old farm implements and 'recreations' of farming processes, together with information displays. This had always been a personal interest of Jim's and it allowed him to convert two existing farm outbuilding to create a 'Museum of Farming Heritage'.

The year after, they introduced tractor rides for visitors around the whole farm and extended the Petting Zoo and farming tradition exhibits further. But the most significant investment was in the 'Preserving Kitchen'. *'We had been looking for some way of using the surplus fruits and vegetable that we occasionally accumulated and also for some kind of products that we could sell in a farm shop. We started the Preserving Kitchen to make jams and fruit, vegetables and sauces preserved in jars. The venture was an immediate success. We started making just 50 kilograms of preserves a week, within 3 months that had grown to 300 kilogrammes a week and we are now producing around 1000 kilogrammes a week, all under the 'Blackberry Hill Farm' label.'* The following year the preserving kitchen was extended and a viewing area added. *'It was a great attraction from the beginning'*, said Mandy, *'We employed ladies* *from the local village to make the preserves. They are all extrovert characters, so when we asked them to dress up in traditional 'farmers wives' type clothing they were happy to do it. The visitors love it, especially the good natured repartee with our ladies. The ladies also enjoy giving informal history lessons when we get school parties visiting us.'*

Within the last two years, the farm had further extended its Preserving Kitchen, farm shop, exhibits and Petting Zoo. It had also introduced a small adventure playground for the children, a café serving drinks and its own produce, a picnic area and a small bakery. The bakery was also open for view by customers and staffed by bakers in traditional dress. *'It's a nice little visitor attraction'*, said Mandy, *'and it gives us another opportunity to squeeze more value out of our own products'*. Table 8.2(a) shows last year's visitor numbers; Table 8.2(b) shows the farm's opening times.

Demand

The number of visitors to the farm is extremely seasonal. From a low point in January and February, when most people just visited the farm shop, the spring and summer months could be very busy, especially on public holidays. The previous year Mandy had tracked the number of visitors arriving at the farm each day. *'It is easy to record the number of people visiting the farm attractions, because they pay the entrance charge. What we had not done before is include the people who just visited the farm shop and bakery that can be accessed both from within the farm and from the car park. We estimate that the number of people visiting the shop but not the farm ranges from 74 per cent in February down to around 15 per cent in August.'* Figure 8.13 shows the number of visitors in August of the previous year. *'What our figures do not include are those people who visit the shop but don't buy anything. This is unlikely to be a large number.'*

Table 8.2(a) Number of visitors last year

Month	*Total visitors*
January	1,006
February	971
March	2,874
April	6,622
May	8,905
June	12,304
July	14,484
August	15,023
September	12,938
October	6,687
November	2,505
December	3,777
Total	88,096
Average	7,341.33

Table 8.2(b) Farm opening times*

January to Mid-March	Wednesday–Sunday	10.00–16.00
Mid-March to May	Tuesday–Sunday	9.00–18.00
May to September	All week	8.30–19.00
October to November	Tuesday–Sunday	10.00–16.00
December	Tuesday–Sunday	9.00–18.00

*Special Evening events Easter, summer weekends and Christmas.

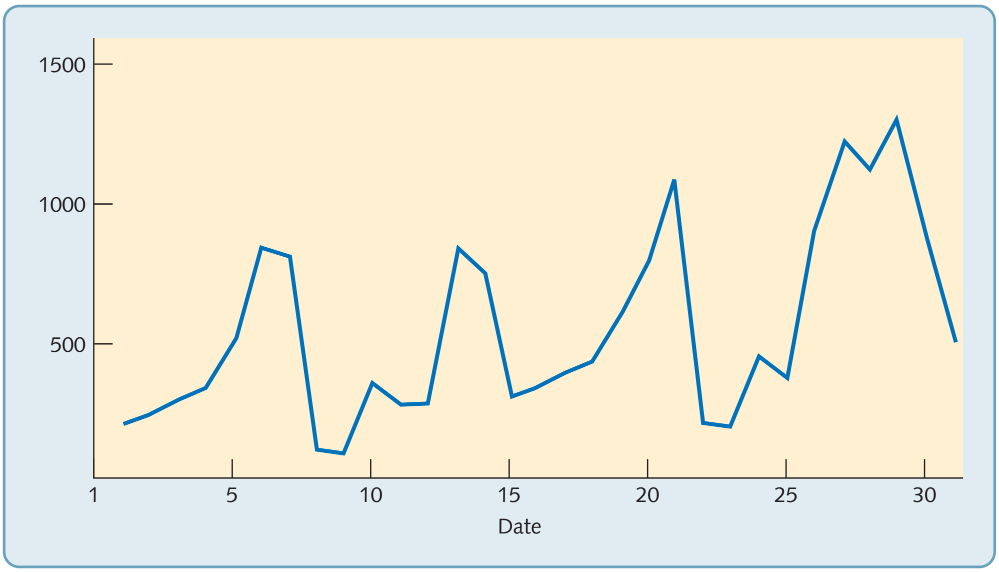

Figure 8.13 Daily number of visitors in August last year

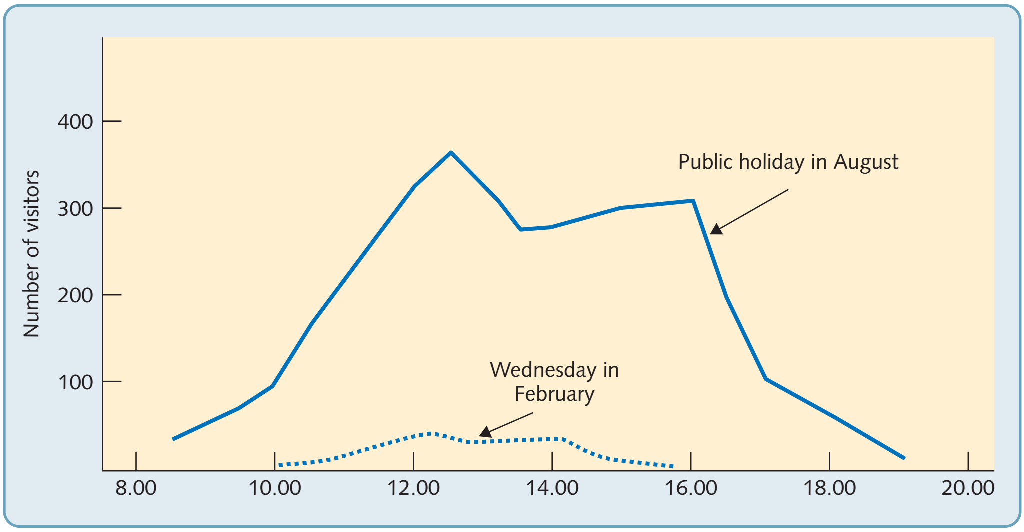

Figure 8.14 Visitor arrivals – public holiday in August and a Wednesday in February

Mandy had also estimated the average stay at the farm and/or farm shop. She reckoned that in winter time the average stay was 45 minutes, but in August in climbed to 3.1 hours.

Current issues

Both Jim and Mandy agreed that their lives had fundamentally changed over the last few years. Income from visitors and from the Blackberry Hill brand of preserves, now accounted for 70 per cent of the farm's revenue. More importantly, the whole enterprise was significantly more profitable than it had ever been. Nevertheless, the farm faced a number of issues.

The first was the balance between its different activities. Jim was particularly concerned that the business remained a genuine farm. *'When you look at the revenue per hectare, visitor and production activities bring in far more revenue than conventional agricultural activities. However, if we push the agri-tourism too far we become no better than a theme park. We represent something more than this to our visitors. They come to us partly because of what we represent as well as what we actually do. I am not sure that we would want to grow much more. Anyway, more visitors would mean that we have to extend the car park. That would be expensive, and although it would be necessary, it does not directly bring in any more revenue. There are already parking problems during peak period and we have had complaints from the police that our visitors park inappropriately on local roads.'*

'There is also the problem of complexity. Every time we introduce a new attraction, the whole business gets that little bit more complex to manage. Although we enjoy it tremendously, both Mandy and I are spreading ourselves thinly over an ever-widening range of activities.' Mandy was also concerned over this. *'I'm starting to feel that my time is being taken up in managing the day-to-day problems of the business. This does not leave time either for thinking about the overall direction in which we should be going, or spending time talking with the staff. That is why we both see this coming year as a time for consolidation and for smoothing out the day-to-day problems of managing the business, particularly the queuing, which is getting excessive at busy times. That is why this year we are limiting ourselves to just one new venture for the business.'*

Staff management was also a concern for Mandy. The business had grown to over 80 (almost all part-time and seasonal) employees. *'We have become a significant employer in the area. Most of our employees are still local people working part-time for extra income, but we are also now employing 20 students during the summer period and, last year, 8 agricultural students from Eastern Europe. But now, labour is short in this part of the country and it is becoming more difficult to attract local people, especially to produce Blackberry Hill Farm Preserves. Half of the Preserving Kitchen staff work all year, with the other employed during the summer and autumn periods. But most of them would prefer guaranteed employment throughout the year.'*

Table 8.3 gives more details of some of the issues of managing the facilities at the farm.

Where next?

By the 'consolidation' and improvement of 'day-to-day' activities, Jim and Mandy meant that they wanted to increase their revenue, while at the same time reducing the occasional queues that they knew could irritate their visitors, preferably without any significant investment in extra capacity. They also were concerned to be able to offer more stable employment to the Preserving Kitchen employees throughout the year, who would produce at a

Table 8.3 The farm's main facilities and some of the issues concerned with managing them

Facility	*Issues*
Car park	• 85 car parking spaces, 4 × 40 -seater tour bus spaces
Fixed exhibits, etc. Recreation of old farmhouse kitchen, recreation of barnyard, old fashioned milking parlour, various small exhibits on farming past and present, adventure playground, ice cream and snack stands	• Most exhibits in, or adjacent to the farm museum • At peak times helpers dressed in period costume entertain visitors • Feedback indicates customers find exhibits more interesting than they thought they would • Visitor free to look when they wish absorbs demand from busy facilities
Tractor rides One tractor tows decorated covered cart with maximum capacity of 30 people, tour takes around 20 minutes on average (including stops). Waits 10 minutes between tours except at peak times when tractor circulates continuously	• Tractor acts both as transport and entertainment, approximately 60 per cent of visitors stay on for the whole tour, 40 per cent use it as 'hop-on hop-off' facility • Overloaded at peak times, long queues building • Feedback indicates it is popular, except for queuing • Jim reluctant to invest in further cart and tractor
Pick-your-own area Largest single facility on the farm. Use local press, dedicated telephone line (answering machine) and website to communicate availability of fruit and vegetables. Checkout and weighing area next to farm shop, also displays picked produce and preserves etc. for sale	• Very seasonal and weather dependent, both for supply and demand • Farm plans for a surplus over visitor demand, uses surplus in preserves • Six weighing/paying stations at undercover checkout area. Queues develop at peak times. Feedback indicates some dissatisfaction with this • Can move staff from farm shop to help with checkout in busy periods, but farm shop also tends to be busy at the same time • Considering using packers at pay stations to speed up the process
Petting Zoo Accommodation for smaller animals including sheep and pigs. Large animals (cattle, horses) brought to viewing area daily. Visitors can view all animals and handle/stroke most animals under supervision	• Approximately 50 per cent of visitors view Petting Zoo. • Number of staff in attendance varies between 0 (off-peak) and 5 (peak periods) • The area can get congested during peak periods • Staff need to be skilled at managing children
Preserving Kitchen Boiling vats, mixing vats, jar sterilising equipment, etc. Visitor viewing area can hold 15 people comfortably. Average length of stay 7 minutes in off-season, 14 minutes in peak season	• Capacity of kitchen is theoretically 4,500 kilogrammes per month on a 5-day week and 6,000 kilogrammes on a 7-day week. • In practice, capacity varies with season because of interaction with visitors. Can be as low as 5,000 kilogrammes on a 7-day week in summer, or up to 5,000 kilogrammes on a 5-day week in winter • Shelf-life of products is on average 12 months • Current storage area can hold 16,000 kilogrammes
Bakery Contains mixing and shaping equipment, commercial oven, cooling racks, display stand, etc. Just installed doughnut making machine. All pastries contain farm's preserved fruit	• Starting to become a bottleneck since doughnut making machine installed; visitors like watching it • Products also on sale at farm shop adjacent to bakery • Would be difficult to expand this area because of building constraints
Farm shop and café Started by selling farm's own products exclusively. Now sells a range of products from farms in the region and wider. Started selling frozen menu dishes (lasagne, goulash, etc.) produced off-peak in the Preserving Kitchen	• The most profitable part of the whole enterprise, Jim and Mandy would like to extend the retailing and café operation • Shop includes area for cooking displays, cake decoration, fruit dipping (in chocolate), etc. • Some congestion in shop at peak times but little visitor dissatisfaction • More significant queuing for café in peak periods • Considering allowing customers to place orders before they tour the farm's facilities and collect their purchases later • Retailing more profitable per square metre than café

Table 8.4 Preserve demand and production (previous year)

Month	*Demand (kg)*	*Cumulative demand (kg)*	*Production (kg)*	*Cumulative product (kg)*	*Inventory (kg)*
January	682	682	4,900	4,900	4,218
February	794	1,476	4,620	9,520	8,044
March	1,106	2,582	4,870	14,390	11,808
April	3,444	6,026	5,590	19,980	13,954
May	4,560	10,586	5,840	25,820	15,234
June	6,014	16,600	5,730	31,550	14,950
July	9,870	26,470	5,710	37,260	10,790
August	13,616	40,086	5,910	43,170	3,084
September	5,040	45,126	5,730	48,900	3,774
October	1,993	47,119	1,570*	50,470	3,351
November	2,652	49,771	2,770*	53,240	3,467
December	6,148	55,919	4,560	57,800	1,881
Average demand	4,660			Average inventory	7,880

*Technical problems reduced production level

near constant rate (see Table 8.4). However, they were not sure if this could be done without storing the products for so long that their shelf-life would be seriously affected. There was no problem with the supply of produce to keep production level; less than 2 per cent of the fruit and vegetables that go into their preserves are actually grown on the farm. The remainder were bought at wholesale markets, although this was not generally understood by customers.

Of the many ideas being discussed for next year's 'one new venture', two were emerging as particularly attractive. Jim liked the idea of developing a Maize Maze, a type of attraction that had become increasingly popular in Europe and North America in the last five years. It involves planting a field of maize (corn) and, once grown, cutting through a complex serious of paths in the form of a maze. Evidence from other farms indicates that a maze would be extremely attractive to visitors and Jim reckons that it could account for up to an extra 10,000 visitors during the summer period. Designed as a separate activity with its own admission charge, it would require an investment of around £20,000, but generate more than twice that in admission charges as well as attracting more visitors to the farm itself.

Mandy favours the alterative idea – that of building up their business in organised school visits. '*Last year we joined the National Association of Farms for Schools. Their advice is that we could easily become one of the top school attractions in this part of England. Educating visitors about farming tradition is already a major part of what we do. And many of our staff have developed the skills to communicate to children exactly what farm life used to be like. We would need to convert and extend one of our existing underused farm outbuildings to make a 'school room' and that would cost between and £30,000 and £35,000. And although we would need to discount our admission charge substantially, I think we could break even on the investment within around two years.*'

QUESTIONS

1 How could the farm's day-to-day operations be improved?

2 What advice would you give Jim and Mandy regarding this year's 'new venture'?

APPLYING THE PRINCIPLES

Some of these exercises can be answered by reading the chapter. Others will require some general knowledge of business activity and some might require an element of investigation. **Model answers for the first two questions are available on the student companion website.** For model answers to the other questions in this section, please ask your tutor/lecturer.

1 A Pizza Company has a demand forecast for the next 12 months that is shown in Table 8.5. The current workforce of 100 staff can produce 1500 cases of pizzas per month.

(a) Prepare a production plan that keeps the output level. How much warehouse space would the company need for this plan?

(b) Prepare a demand chase plan. What implications would this have for staffing levels, assuming that the maximum amount of overtime would result in production levels of only 10 per cent greater than normal working hours?

Table 8.5 Pizza demand forecast

Month	*Demand (cases per month)*
January	600
February	800
March	1,000
April	1,500
May	2,000
June	1,700
July	1,200
August	1,100
September	900
October	2,500
November	3,200
December	900

2 In a typical 7-day period, the planning department of the pizza company programs its 'Pizzamatic' machine for 148 hours. It knows that changeovers and set-ups take 8 hours and breakdowns average 4 hours each week. Waiting for ingredients to be delivered usually accounts for 6 hours, during which the machine cannot work. When the machine is running, it averages 87 per cent of its design speed. And inspection has revealed that 2 per cent of the pizzas processed by the machine are not up to the company's quality standard. Calculate the OEE of the 'Pizzamatic' machine.

3 A German car manufacturer defines 'utilisation' as the ratio of actual output for a process to its design capacity, where design capacity is the capacity of a process as it is designed to operate. However, it knows that it is rarely possible to achieve this theoretical level of capacity. This is why the company uses a measure that it calls 'effective capacity', which is the actual capacity of a process, once maintenance, changeover, other stoppages, and loading have been considered. The ratio of actual output for a process to its effective capacity is defined as its 'efficiency'.

The company has a painting line with a design capacity of 100 m^2 per minute and the line is operated 24 hours per day, 7 days a week (168 hours). Records for a week show lost time in production as given in Table 8.6.

During this week, production was only $100 \times 60 \times (168 - 100) = 408{,}000$ m^2 per week. What is the painting line's 'utilisation' and 'efficiency' according to the company's definitions?

Table 8.6 Lost time in production for German car manufacturer

	Activity	*Number of hours lost (h)*
1	Product changeovers (set-ups)	18
2	Regular maintenance	12
3	No work scheduled	6
4	Quality sampling checks	8
5	Shift change times	8
6	Maintenance breakdown	16
7	Quality failure investigation	12
8	Paint stock-outs	6
9	Labour shortages	6
10	Waiting for paint	5
	Total	100

4 Seasonal demand is particularly important to the greetings card industry. Mother's Day, Father's Day Halloween, Valentine's Day and other occasions have all been promoted as times to send (and buy) appropriately designed cards. Now, some card manufacturers have moved on to 'non-occasion' cards, which can be sent at any time. The cards include those intended to be sent from a parent to a child with messages such as 'Would a hug help?', 'Sorry I made you feel bad', and 'You're perfectly wonderful – it's your room that's a mess'. Other cards deal with more serious adult themes such as friendship ('you're more than a friend, you're just like family') or even alcoholism ('this is hard to say, but I think you're a much neater person when you're not drinking'). Some card companies have founded 'loyalty marketing groups' that 'help companies communicate with their customers at an emotional level'. They promote the use of greetings cards for corporate use, to show that customers and employees are valued

(a) What seem to be the advantages and disadvantages of these strategies?

(b) What else could card companies do to cope with demand fluctuations?

5 Revisit the example, 'United drags passengers off its plane'. (a) How should the airline have handled the situation? (b) After the incident attracted so much negative publicity, United announced a new upper limit of $10,000 in compensation for passengers who agree to give up a seat on a flight where United needs to free space and that it would create a 'customer solutions team to provide agents with creative solutions' for getting inconvenienced customers to their destination. Do you think that these were sensible moves? (c) Within a few days another 'scandal' hit the airline. A 'potentially prize-winning' rabbit (called Simon) reportedly died while in transit from London Heathrow to O'Hare airport in Chicago. Why is this incident so important to United?

Notes on chapter

1 Sources include: *The Economist* (2011) 'Amazon: The Wal-Mart of the web, 1 October; Amazon website (2014) www.amazon.co.uk; Reuters.com (2013) 'Amazon plans big expansion of online grocery business', 4 June; Ralph, A. (2011) 'Too many unhappy returns', *The Times,* 26 November.

2 Sources include: *The Economist* (2009) 'A piece of cake: Panettone season arrives', *The Economist* print edition, 10 December; Bauli website, http://www.bauligroup.it/en/

3 Sources include: Equants (2012) 'London Heathrow airport – Economic and social impact', Ecquants.com/2012_LHR.aspx; Heathrow Airport Holdings Company (2014) Heathrow Airport Holdings Company Information, Retrieved 12 October 2014.

4 Sources include: Horticulture Week (2010) 'Grower guide tackles problem of building sustainable workforces', Friday, 14 May; Growing Jobs website, http://growingjobs.org/

5 Source: The University of Limerick, (2015) 'A study of the prevalence of Zero Hours Contracts'.

6 Sources include: *The Economist* (2016) 'Jacking up prices may not be the only way to balance supply and demand for taxis', Free exchange, 14 May; Dholakia, U. M. (2015) Everyone hates Uber's surge pricing – here's how to fix it', *Harvard Business Review,* 21 December. Note – information is correct at the time of writing (early 2017) please check to see if the company have changed their policy.

7 Kimes, S. (1989) 'Yield management: a tool for capacity-constrained service firms', *Journal of Operations Management,* vol. 8, no. 4.

8 Sources include: Paton, G. (2017) 'Airline price war means you could be dragged off your next flight', *The Times,* 11 April; BBC News site (2017) 'United CEO says removed passenger was "disruptive and belligerent"', 11 April; Gunter, J. (2017) 'United Airlines incident: Why do airlines overbook?', *BBC News,* 10 April; Hill, A. (2017) 'United's reputational repair job mixes sense and nonsense', *Financial Times,* 27 April.

TAKING IT FURTHER

Gunther, Neil, J. (2007) *Guerrilla Capacity Planning, Springer.* *This book provides a tactical approach for planning capacity in both product-based and service-based contexts. Particularly interesting for those new to the ideas of capacity planning as it covers basic and more advanced demand forecasting techniques as well as 'classic' capacity responses.*

Hansen, Robert, C. (2005) *Overall Equipment Effectiveness (OEE), Industrial Press.* *If you want to know more about OEE, its origins and application, this is the place to start.*

Manas, J. (2014) *The Resource Management and Capacity Planning Handbook: A Guide to Maximizing the Value of Your Limited People Resources, McGraw-Hill Education.* *Very much a practitioners guide.*

Van Mieghem, J. (2003) *'Capacity Management, Investment, and Hedging: Review and Recent Developments', Manufacturing and Service Operations Management, vol. 5, no. 4.* *An academic article reviewing the literature on strategic capacity management. It does a nice job of covering the different approaches to capacity management under conditions of stability versus volatility (demand change) and of certainty versus uncertainty i.e. (the predictability of change).*

9 Inventory management

Introduction

Operations managers often have an ambivalent attitude towards inventories – of material, of people (queues), or of information. They can be costly, tying up working capital, they can annoy customers by making them wait, they are risky because items or information held in stock could deteriorate, become obsolete or just get lost, and they can take up valuable space in the operation. On the other hand, inventories can provide some security in an uncertain environment. Knowing that you have the things, people or information 'in stock' is a comforting insurance against demand fluctuations. This is the dilemma of inventory management: in spite of the cost and the other disadvantages associated with holding stocks, they do facilitate the smoothing of supply and demand. In fact, they only exist because supply and demand are not exactly in harmony with each other. Figure 9.1 shows the position of the ideas described in this chapter in the general model of operations management.

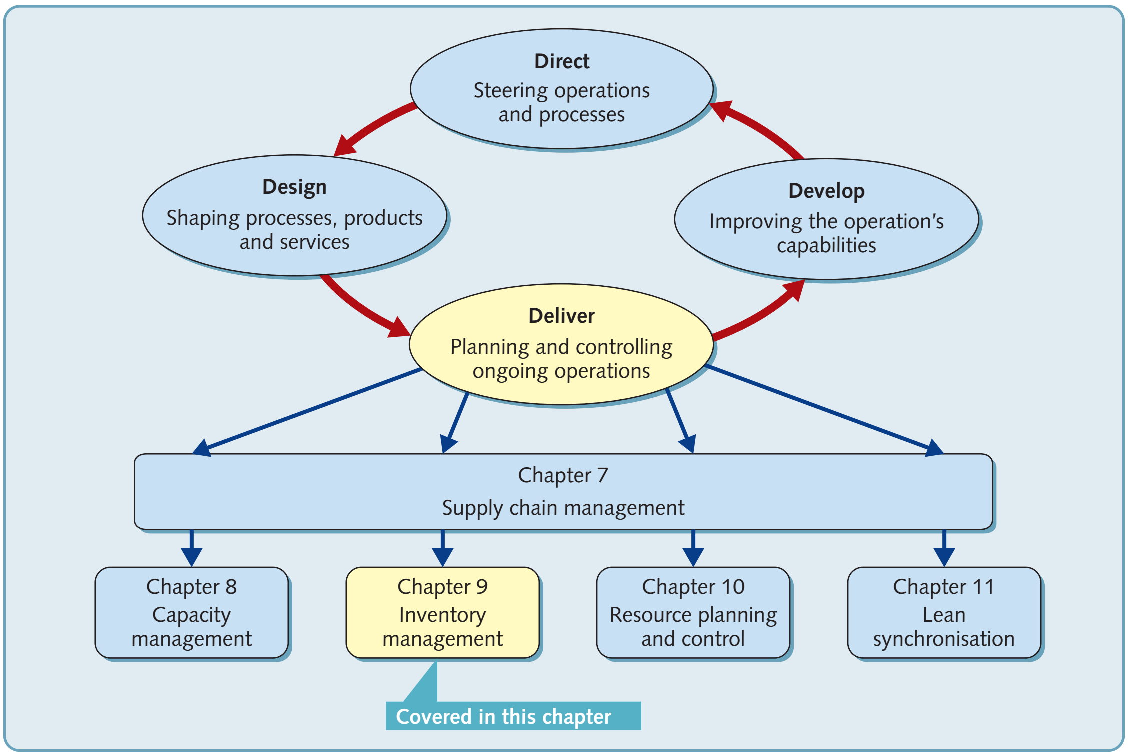

Figure 9.1 Inventory management is the activity of planning and controlling accumulations of transformed resources as they move through supply networks, operations and processes

EXECUTIVE SUMMARY

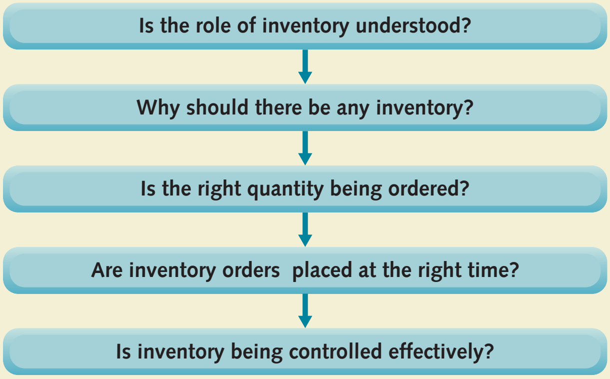

Is the role of inventory understood?

Inventory management is the activity of planning and controlling accumulations of transformed resources as they move through supply networks, operations and processes. The inventory can be accumulations of materials, customers or information. Accumulations of inventory occur because of local mismatches between supplier and demand. All operations have inventories of some kind and inventory management is particularly important where the inventories are central to the operation's objectives and/or of high value. How inventories are managed will determine the balance between customer service and cost objectives.

Why should there be any inventory?

Although there are cost, space, quality and operational / organisational disadvantages with inventory, there are also benefits. As far as physical inventory is concerned, it can act as an insurance against uncertainty, or a buffer against unexpected fluctuations in supply and demand, or to counteract a lack of flexibility, or to allow operations to take advantage of short-term opportunities, to anticipate future demands, to reduce overall costs, or even because it may increase in value. Inventories (queues) of customers can again help balance capacity and demand, or can enable prioritisation, or give customers time to choose, and enable efficient use of resources. Inventories of information (databases) can provide efficient multi-level access, allow single data capture and speed some processes. The underlying objective of inventory management is often to minimise inventory while maintaining acceptable customer service.

Is the right quantity being ordered?

A key inventory decision is the 'order-quantity' decision. Various formulae exist that attempt to identify the order quantity that minimises total costs under different circumstances. One approach to this problem, the newsvendor problem, includes the effects of probabilistic demand in determining order quantity.

Are inventory orders placed at the right time?

Broadly, there are two approaches to this. The reorder point approach is to time reordering at the point in time where stock will fall to zero minus the order lead time. A variation of this is to reorder at the equivalent inventory level (the reorder level approach). Reordering at a fixed point or level is termed continuous review methods because it requires continuous monitoring of stock levels. A different approach, called the periodic review approach, places orders at predetermined times, but varies the order depending on the level of inventory at that time. Both continuous and periodic review can be calculated on a probabilistic basis to include safety stocks.

Is inventory being controlled effectively?

The most common inventory control approach is based on the Pareto (at 80:20) curve. It classifies stocked items by the usage value (their usage rate multiplied by their value). High-usage value items are deemed A class and controlled carefully, whereas low-usage value items (B and C class) are controlled less intensely. However, this approach often has to be modified to take account of slow-moving items. Inventory information systems are generally used to keep track of inventory, forecast demand and place orders automatically.

DIAGNOSTIC QUESTION

Is the role of inventory understood?

Inventory is a term we use to describe the accumulations of materials, customers or information as they flow through processes or networks. Occasionally, the term is also used to describe transforming resources, such as rooms in hotels or automobiles in a vehicle hire firm, but here we use the term for the accumulation of resources that flow through processes, operations or supply networks. Physical inventory (sometimes called 'stock') is the accumulation of physical materials such as components, parts, finished goods or physical (paper) information records. Queues are accumulations of customers, physical as in a queuing line or people in an airport departure lounge, or waiting for service at the end of phone lines. Databases are stores for accumulations of digital information, such as medical records or insurance details. Managing these accumulations is what we call 'inventory management'. And it's important. Material inventories in a factory can represent a substantial proportion of cash tied up in working capital. Minimising them can release large quantities of cash. However, reducing them too far can lead to customers' orders not being fulfilled. Customers held up in queues for too long can get irritated, angry and possibly leave, so reducing revenue. Databases are critical for storing digital information and while storage may be inexpensive, maintaining databases may not be.

OPERATIONS PRINCIPLE

All processes, operations and supply networks have inventories (accumulations) of materials, customers and information.

Inventories are the result of uneven flow

Most things that flow do so in an uneven way. Rivers flow faster down steep sections or where they are squeezed into a ravine. Over relatively level ground, they flow slowly, and form pools or even large lakes where there are natural or man-made barriers blocking its path. It's the same in operations. Passengers in an airport flow from public transport or their vehicles, then have to queue at several points including check-in, security screening and immigration. They then have to wait (a queue even if they are sitting) again in the departure lounge as they are joined (batched) with other passengers to form a group of several hundred people who are ready to board the aircraft. They are then squeezed down the air bridge as they file in one at a time to board the plane. Likewise, in a tractor assembly plant, stocks of components such as gearboxes, wheels, lighting circuits, and so on, are brought into the factory in 10s or 100s and are then stored next to the assembly line ready for use. Finished tractors will also be stored until the transporter comes to take them away in ones or 10s to the dealers or directly to the end customer. Similarly, a government tax department collects information about us and our finances from various sources, including our employers, our tax forms, information from banks or other investment companies, and stores this in databases until they are checked, sometimes by people, sometimes automatically, to create our tax codes and/or tax bills. In fact, because most operations involve flows of materials, customers and/or information, at some points they are likely to have material and information inventories and queues of customers waiting for goods or services, see Table 9.1.

Inventories are often the result of uneven flows. If there is a difference between the timing of supply and demand, or the rate of supply and demand at any point in a process or network then accumulations will occur. A common analogy is the water tank shown in Figure 9.2. If, over time, the rate of supply of water to the tank differs from the rate at which it is demanded, a tank of water (inventory) will be needed to maintain supply. When the rate of supply exceeds the rate

Table 9.1 Examples of inventory held in processes, operations or supply networks

Process, operation or supply network	*Inventories*		
	Physical inventories	*Queues of customers*	*Information in databases*
Hotel	Food items, drinks, toilet items	At check-in and check-out	Customer details, loyalty card holders, catering suppliers
Hospital	Dressings, disposable instruments, blood	Patients on a waiting list, patients in bed waiting for surgery, patients in recovery wards	Patient medical records
Credit card application process	Blank cards, form letters	Customers waiting on the phone	Customers' credit and personal information
Computer manufacturer	Components for assembly, packaging materials, finished computers ready for sale	Customers waiting for delivery of their computer	Customers' details, supplier information

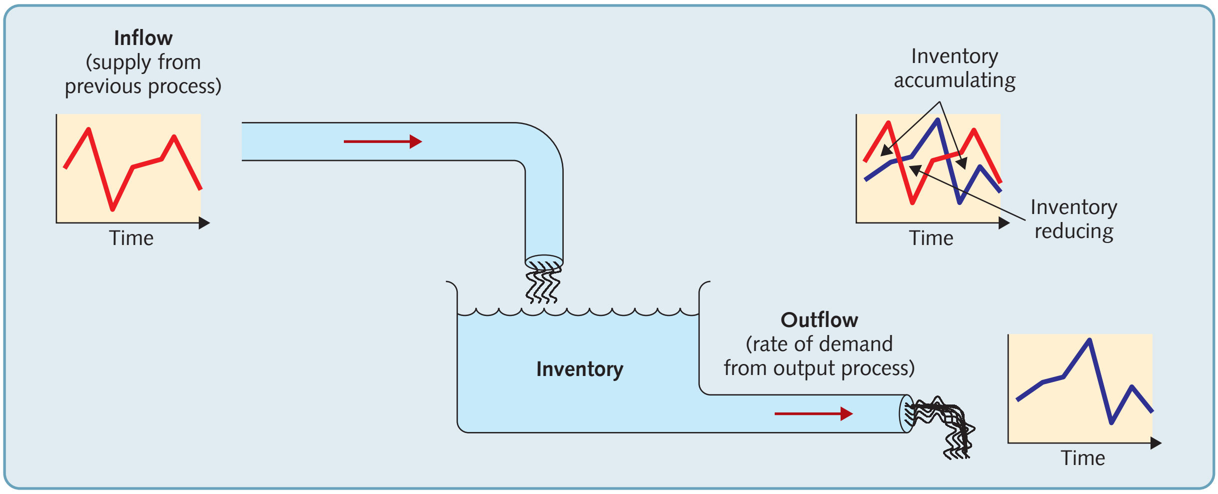

Figure 9.2 Inventory is created to compensate for the differences in timing between supply and demand

of demand, inventory increases; when the rate of demand exceeds the rate of supply, inventory decreases. So if an operation or process can match supply and demand rates, it will also succeed in reducing its inventory levels. But most organisations must cope with unequal supply and demand, at least at some points in their supply chain. Both the following organisations depend on the ability to manage supply and demand inequality through their inventory management.

There is a complication when using this 'water flow' analogy to represent flows and accumulations (inventories) of information. Inventories of information can either be stored because of uneven flow, in the same way as materials and people, or stored because the operation needs to use the information to process something in the future. For example, an internet retail operation will process each order it receives, and inventories of information may accumulate because of uneven flows as we have described. In addition, during order processing customer details could be permanently stored in a database. This information will then be used, not only for future orders from the same customer, but also for other processes, such as targeting promotional activities. In this case, the inventory of information had turned from a transformed resource into a transforming resource, because it is being used to transform other information rather than being transformed itself. So, whereas managing physical material concerns ordering and holding the right amounts of goods or materials to deal with the variations in flow, and managing queues is about the level of resources to deal with demand, a database is the accumulation of information but may not cause an interruption to the flow. Managing databases is about the organisation of the data, its storage, security and retrieval (access and search).

Here are two examples that demonstrate the importance of managing inventory.

EXAMPLE

The blood bank's 'perfect storm[1]

Inventory depends on both supply and demand, so when both are uncertain, inventory management poses particular challenges. When, in addition, the consequences of running out of stock can affect people's health, and then inventory management becomes a particularly vital task. Welcome to the world of the Blood Stocks Management Scheme of the National Health Service Blood and Transplant (NHSBT) that manages blood stocks across the blood supply chain in the UK. NHSBT is responsible for the collection, processing, testing and issuing of blood across England and North Wales. Each year approximately 2 million blood donations are collected from 1.4 million donors to supply hospitals with all the blood needed for accident and emergency situations and regular medical treatment. Many people owe their lives to transfusions that were made possible by the efficient management of blood, stocked in a supply network that stretches from donation centres through to hospital blood banks. The blood supply chain has three main stages.

1. Collection, which involves recruiting and retaining blood donors, encouraging them to attend donor sessions and transporting the donated blood.
2. Processing, that breaks blood down into its constituent parts.
3. Distribution, that transports blood from blood centres to hospitals in response to both routine and emergency requests.

Inventory accumulates at all three stages, as well as in individual hospitals' blood banks. Within the supply chain, less than 10 per cent of donated red blood cells are lost. Much of this is due to losses in processing, but around 5 per cent is not used because it has 'become unavailable', mainly because it has been stored for too long. Part of the inventory management task is to keep this 'time expired' loss to a minimum. In fact, most blood is lost when it is stored in hospital blood banks that are outside the service's direct control. Additionally, blood components will deteriorate over time. Platelets have a shelf life of only 5 days and demand can fluctuate significantly, which makes stock control particularly difficult. Even red blood cells that have a shelf life of 35 days may not be acceptable to hospitals if they are close to their 'used by date'. Stock accuracy is crucial. Giving a patient the wrong type of blood can be fatal.

At a local level, accidents can significantly affect demand. One serious accident involving a cyclist used 750 units of blood, which completely exhausted the available supply (miraculously, he survived). Large-scale accidents usually generate a surge of offers from donors wishing to make immediate donations. There is also a more predictable seasonality to donating blood, however, with a low period during the summer vacation. During public holidays and sporting events, blood donations drop. For example, on one day when the football World Cup quarter final, and Andy Murray's (a British tennis player) Wimbledon semi-final coincided, there was a 12 per cent drop in donations compared with the previous year. The summer of 2012 proved particularly difficult with a cluster of events and public holidays between April and August, including the Queen's Jubilee, Euro 2012, the London Olympic Games and the Paralympics Games. Not only did these events reduce donations (supply), the increased number of visitors to London increased demand. Before the period NHS Blood and Transplant said that the number of major events would create a ‘perfect storm’ and dramatically impact the number of blood donations coming in.

EXAMPLE

Mountains of grit[2]

Students of operations management from Singapore to Saudi Arabia will maybe not have a full appreciation of how important this decision is in the colder parts of the world, but, believe me, road gritting is big news every winter where snow and ice can cause huge disruption to everyday life. But not every time it snows, and more interestingly, not everywhere it snows. The local government authorities around northern Europe and America differ significantly in how well they cope with freezing weather, usually by spreading grit (actually rock salt, a mixture of salt and grit) on the roads. So how do the authorities decide how much grit to stock up in preparation for winter, and when to spread it on the roads? For example, in the UK, when snow is forecast, potential trouble spots are identified by networks of sensors embedded in the road surface to measure climatic conditions. Each sensor is connected by cable or mobile phone technology to an automatic weather station by the roadside. The siting of the sensors is important. They must be sited either on a representative stretch of road (no nearby trees, buildings or bridges, which offer some protection from the cold), or traditional cold spots. The weather stations then beam back data about air and road temperatures, wind speed and direction and the wetness of roads. Salt levels are also measured to ensure that grit that has already been spread has not been blown away by wind or washed away by rain. It has been known for cold weather to be forecast and the gritting trucks to be dispatched, only for the weather to change, with snow turning to rain, which washed away the grit. Then when temperatures suddenly drop again the rain freezes on the road. However, forecasting how much grit will be needed is even more difficult. Long-range weather forecasts are notoriously inaccurate, so no one knows just how bad a coming winter will be. To make matters worse, the need for road grit depends on more than just the total volume of snow. Local authorities can use the same amount of salt on one 30 cm snowfall as one 5 cm snowfall. Furthermore, the number of snowy days is important in determining how much grit will be needed. In the skiing areas of central Europe, most winter days will have snow predictably, while parts of the UK could have little or no snow one winter and many weeks of snow the next.

Supplies of road grit can also vary, as can its price. There are many reasons for this. Mainly of course, if a bad winter is forecast, all authorities in an area will want to buy the same grit, which will reduce supply and put prices up. Furthermore, salt mines can flood, especially in winter. Nor is it cheap to transport grit from one area to another; it is a low value but heavy material. Consequently, some authorities organise purchasing groups to get better prices before the season starts. Getting more salt during the season may be possible but prices are higher and supply is not guaranteed. In addition, an authority has do decide how fast to use up its inventory of grit. At the start of the winter period, authorities may be cautious about gritting because, once used, the grit cannot be used again, and who knows what the weather will be like later in the season. But in the final analysis, the decision of how large an inventory of grit to buy and how to use it is a balance between risks and consequences. If too big an inventory of grit is built up it may not all be used, with the cost of carrying it over to next year being borne by local taxpayers. Too small an inventory may incur the wrath of local voters when the roads are difficult to negotiate. Of course, a perfect weather forecast would help!

What do they have in common?

Both of these organisations depend on their ability to manage inventory. In doing so, both are attempting to manage the trade-off that lies at the heart of all inventory management; balancing the costs of holding stock against the customer service that comes from having appropriate stock levels. Stock levels that are too high have a cost. This may be simply working capital in the case of road grit, or it could be the cost of blood becoming out-of-date and wasted in the blood service. Without an appropriate level of inventory, customers suffer poor service. This means potentially disrupting local traffic flow in the case of road gritting. But a failure of supply for the blood service may have even more drastic consequences. For both operations at each point in the inventory system, operations managers need to manage the day-to-day tasks of running the system. Requests for grit or blood will be received from internal or external customers; they will be supplied and demand will gradually deplete the inventory. Orders will need to be placed for replenishment of the stocks; deliveries will arrive and require storing. At each stage of managing the inventory, the competing demands of costs and service levels will need trading-off. In order to manage this trade-off we first need to understand the reasons for not having inventory, the reasons for having it, then understand the tools available to help make these balancing decisions.

DIAGNOSTIC QUESTION

Why should there be any inventory?

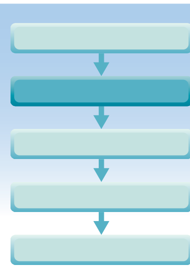

There are plenty of reasons to avoid accumulating inventory where possible. Table 9.2 identifies some of these, particularly those concerned with cost, space, quality and operational / organisational issues.

Table 9.2 Some reasons to avoid inventories

	Inventories		
	Physical inventories	*Queues of customers*	*Digital information in databases*
Cost	Ties up working capital and there could be high administrative and insurance costs	Primarily time-cost to the customer, i.e. wastes customers' time	Cost of set-up, access, update and maintenance
Space	Requires storage space	Requires areas for waiting or phone lines for held calls	Requires memory capacity. May require secure and/or special environment
Quality	May deteriorate over time, become damaged or obsolete	May upset customers if they have to wait too long. May lose customers	Data may be corrupted or lost or become obsolete
Operational / organisational	May hide problems (see lean synchronisation, Chapter 11)	May put undue pressure on the staff and so quality is compromised for throughput	Databases need constant management; access control, updating and security

So why have inventory?

On the face of it, it may seem sensible to have a smooth and even flow of materials, customers and information through operational processes and networks and thus not have any accumulations. In fact, inventories provide many advantages for both operations and their customers. If a customer has to go to a competitor because a part is out of stock, or because they have had to wait too long, or because the company insists on collecting all their personal details each time they call, the value of inventories seems undisputable. The task of operations management is to allow inventory to accumulate only when its benefits outweigh its disadvantages. The following are some of the benefits of inventory.

OPERATIONS PRINCIPLE

Inventory should only accumulate when the advantages of having it outweigh its disadvantages.

Physical inventory is an insurance against uncertainty

Inventory can act as a buffer against unexpected fluctuations in supply and demand. For example, a retail operation can never forecast demand perfectly over the lead time. It will order goods from its suppliers such that there is always a minimum level of inventory to cover against the possibility that demand will be greater than expected during the time taken to deliver the goods. This is buffer, or safety inventory. It can also compensate for the uncertainties in the process of the supply of goods into the store. The same applies with the output inventories, which is why hospitals always have a supply of blood, sutures and bandages for immediate response to accident and emergency patients. Similarly, auto-servicing services, factories and airlines may hold selected critical spare parts inventories so that maintenance staff can repair the most common faults without delay. Again, inventory is being used as an 'insurance' against unpredictable events.

Physical inventory can counteract a lack of flexibility

Where a wide range of customer options is offered, unless the operation is perfectly flexible, stock will be needed to ensure supply when it is engaged on other activities. This is sometimes called cycle inventory. For example, suppose a baker makes three types of bread. Because of the nature of the mixing and baking process, only one kind of bread can be produced at any time. The baker will have to produce each type of bread in batches large enough to satisfy the demand for each kind of bread between the times when each batch is ready for sale. So, even when demand is steady and predictable, there will always be some inventory to compensate for the intermittent supply of each type of bread.

Physical inventory allows operations to take advantage of short-term opportunities

Sometimes opportunities arise that necessitate accumulating inventory, even when there is no immediate demand for it. For example, a supplier may offer a particularly good deal on selected items for a limited time period, perhaps because they want to reduce their own finished goods inventories. Under these circumstances, a purchasing department may opportunistically take advantage of the short-term price advantage.

Physical inventory can be used to anticipate future demands

Medium-term capacity management (covered in Chapter 8) may use inventory to cope with demand–capacity. Rather than trying to make a product (such as chocolate) only when it is needed, it is produced throughout the year ahead of demand and put into inventory until it is needed. This type of inventory is called anticipation inventory and is most commonly used when demand fluctuations are large but relatively predictable.

EXAMPLE

Toilet roll delay[3]

Inventories do not just need to be big enough to serve their purpose; they also need to be in the right place. British Airways found this out when a shortage of toilet paper and 'the wrong

kind of headphones' delayed a London to Barbados flight for five hours. The Boeing 777 flight, was due to leave the London airport at 1.40 p.m. on Sunday but was delayed until 6.51 p.m. It eventually arrived at Bridgetown five-and-a-half hours later than scheduled. The problem, apparently, was that a newly appointed supply company had failed to load the right supplies on board, and by the time the supplies arrived, the outbound crew was 'out of hours'. This meant that they could not operate the service within predetermined time limits. It then took three more hours to find another crew. Not only was it embarrassing for the airline, it put them in danger of having to face an almost £300,000 bill for compensation. European rules on flight delays entitle each passenger to compensation of €600. Even worse, the delay to the outbound flight meant that the return service was also delayed by almost six hours, pushing up the compensation bill even further.

Physical inventory can reduce overall costs

Holding relatively large inventories may bring savings that are greater than the cost of holding the inventory. This may occur when bulk-buying gets the lowest possible cost of inputs, or when large order quantities reduce both the number of orders placed and the associated costs of administration and material handling. This is the basis of the 'economic order quantity' (EOQ) approach, which will be discussed later in this chapter.

Physical inventory can increase in value

Sometimes the items held as inventory can increase in value and so become an investment. For example, dealers in fine wines are less reluctant to hold inventory than dealers in wine that does not get better with age. (However, it can be argued that keeping fine wines until they are at their peak is really part of the overall process rather than inventory as such.) A more obvious example is inventories of money. The many financial processes within most organisations will try to maximise the inventory of cash they hold because it is earning them interest.

Physical inventory fills the processing 'pipeline'

'Pipeline' inventory exists because transformed resources cannot be moved instantaneously between the point of supply and the point of demand. When a retail store places an order, its supplier will 'allocate' the stock to the retail store in its own warehouse, pack it, load it onto its truck, transport it to its destination and unload it into the retailer's inventory. From the time that stock is allocated (and therefore it is unavailable to any other customer) to the time it becomes available for the retail store, it is pipeline inventory. Especially in geographically dispersed supply networks, pipeline inventory can be substantial.

Queues of customers help balance capacity and demand

This is especially useful if the main service resource is expensive; for example, doctors, consultants, lawyers or expensive equipment such as CAT scans. By waiting a short time after their arrival, and creating a queue of customers, the service always has customers to process. This is also helpful where arrival times are less predictable, for example where an appointment system is not used or not possible.

Queues of customers enable prioritisation

In cases where resources are fixed and customers are entering the system with different levels of priority, the formation of a queue allows the organisation to serve urgent customers while keeping other less urgent ones waiting. In the UK it is not usual to have to wait for 3 to 4 hours for treatment in an accident and emergency ward, with more urgent cases 'jumping the queue' for treatment.

Queuing gives customers time to choose

Time spent in a queue gives customers time to decide what products/services they require; for example, customers waiting in a fast food restaurant have time to look at the menu so that when they get to the counter they are ready to make their order without holding up the server.

Queues enable efficient use of resources

By allowing queues to form, customers can be batched together to make efficient use of operational resources. For example, a queue for an elevator makes better use of its capacity; in an airport, by calling customers to the gate, staff can load the aircraft more efficiently and quickly.

Databases provide efficient multi-level access

Databases are relatively cheap ways of storing information and providing many people with access, although there may be restrictions or different levels of access. The doctor's receptionist will be able to call up a patient's records to check their name and address, and make an appointment. The doctor will then be able to call up the appointment and the patient's records; and the pharmacist will be able to call up the patients' name and prescriptions and cross-check for other prescriptions and known allergies, and so on.

Databases of information allow single data capture

There is no need to capture data at every transaction with a customer or supplier, though checks may be required.

Databases of information speed the process

Amazon, for example, stores, if you agree, your delivery address and credit card information so that purchases can be made with a single click, making it fast and easy for the customer.

The effect of inventory on return on assets

One can summarise the effects on the financial performance of an operation by looking at how some of the factors of inventory management impact on 'return on assets', a key financial performance measure. Figure 9.3 shows some of these factors.

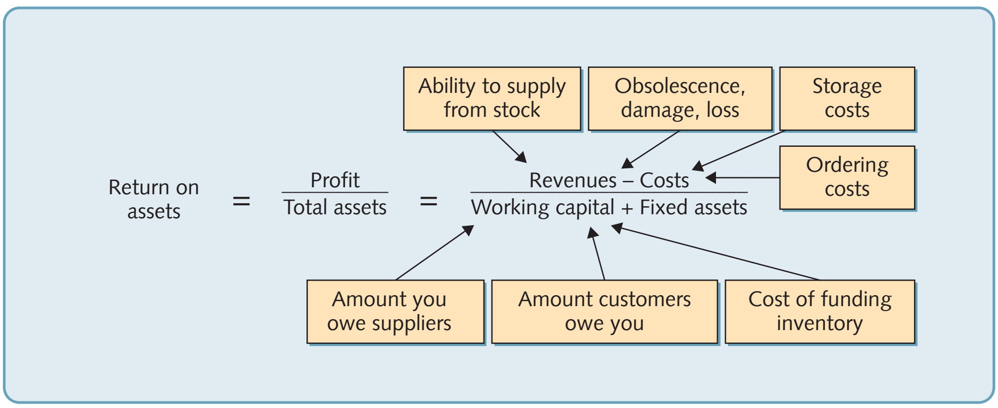

Figure 9.3 **Inventory management has a significant effect on return on assets**

- *Inventory governs the operation's ability to supply its customers.* The absence of inventory means that customers are not satisfied with the possibility of reduced revenue.
- *Inventory may become obsolete as alternatives become available, or could be damaged, deteriorate, or simply get lost.* This increases costs (because resources have been wasted) and reduces revenue (because the obsolete, damaged or lost items cannot be sold*).*
- *Inventory incurs storage costs (leasing space, maintaining appropriate conditions, etc.).* This could be high if items are hazardous to store (e.g., flammable solvents, explosives, chemicals) or difficult to store requiring special facilities (e.g., frozen food).
- *Inventory involves administrative and insurance costs.* Every time a delivery is ordered, time and costs are incurred.
- *Inventory ties up money,* in the form of working capital, which is therefore unavailable for other uses, such as reducing borrowings or making investment in productive fixed assets (we shall expand on the idea of working capital later).
- *Inventory contracts with suppliers can dictate the timing of when suppliers need to be paid.* If they require paying before the operation receives payment from its customers (as is normal), the difference between the amount the operation owes suppliers and the amount suppliers owe the operation, adds to working capital requirements.

OPERATIONS PRINCIPLE

Inventory management can have a significant effect on return on assets.

Reducing physical inventory

The objective of most operations managers who manage physical inventories is to reduce the overall level (and/or cost) of inventory, while maintaining an acceptable level of customer service. Table 9.3 identifies some of the ways in which inventory may be reduced.

Table 9.3 Some ways in which physical inventory may be reduced

Reason for holding inventory	*Example*	*How inventory could be reduced*
As an insurance against uncertainty	Safety stocks for when demand or supply is not perfectly predictable	• Improve demand forecasting • Tighten supply, e.g. through service level penalties
To counteract a lack of flexibility	Cycle stock to maintain supply when other products are being made	• Increase flexibility of processes, e.g. by reducing changeover times (see Chapter 11) • Using parallel processes producing output simultaneously (se Chapter 6)
To take advantage of relatively short-term opportunities	Suppliers offer 'time-limited' special low-cost offers	• Persuade suppliers to adopt 'everyday low prices' (see Chapter 11)
To anticipate future demands	Build up stocks in low-demand periods for use in high-demand periods	• Increase volume flexibility by moving towards a 'chase demand' plan (see Chapter 8)
To reduce overall costs	Purchasing a batch of products in order to save delivery and administration costs	• Reduce administration costs through purchasing process efficiency gains • Investigate alternative delivery channel that reduce transport costs
To fill the processing 'pipeline'	Items being delivered to customer	• Reduce process time between customer request and dispatch of items • Reduce throughput time in the downstream supply chain (see Chapter 7)

DIAGNOSTIC QUESTION

Is the right quantity being ordered?

To illustrate this decision, consider how we manage our domestic inventory. We implicitly make decisions on order quantity, that is, how much to purchase at one time by balancing two sets of costs: the costs associated with going out to purchase the food items and the costs associated with holding the stocks. The option of holding very little or no inventory of food and purchasing each item only when it is needed it requires little money because purchases are made only when needed, but involves buying several times a day, which is inconvenient. Conversely, making one journey to the local superstore every few months and purchasing all the provisions we would need until our next visit reduces purchasing time and costs but requires a very large amount of money each time the trip is made – money that could otherwise be in the bank and earning interest. We might also have to invest in extra cupboard units and a very large freezer. Somewhere between these extremes lies an ordering strategy that will minimise the total costs and effort involved in purchasing food.

Inventory costs

A similar range of costs applies in commercial order-quantity decisions as in the domestic situation. These are costs of placing an order, including preparing the documentation, arranging for the delivery to be made, arranging to pay the supplier for the delivery and the general costs of keeping all the information that allows us to do this. An 'internal order' on processes within an operation has equivalent costs. Price discount costs for large orders or extra costs for small orders may also influence how much to purchase. If inventory cannot supply demand, there will be costs to us incurred by failing to supply customers. External customers may take their business elsewhere. Internal stock-outs could lead to idle time at the next process, inefficiencies and eventually, again, dissatisfied external customers. There are the working capital costs of funding the lag between paying suppliers and receiving payment from customers. Storage costs are the costs associated with physically storing goods, such as renting, heating and lighting a warehouse, as well as insuring the inventory. While stored as inventory there is a risk of obsolescence costs if the inventory is superseded (in the case of a change in fashion) or deteriorates with age (in the case of most foodstuffs).

Some of these costs will decrease as order size is increased; the first three costs (cost of placing an order, price discount costs and stock-out costs) are like this. The other costs (working capital, storage and obsolescence costs) generally increase as order size is increased. But, it may not be the same organisation that incurs each cost. For example, sometimes suppliers agree to hold consignment stock. This means that they deliver large quantities of inventory to their customers to store but will only charge for the goods as and when they are used. In the meantime, they remain the supplier's property so do not have to be financed by the customer, who does however provide storage facilities.

Inventory profiles

An inventory profile is a visual representation of the inventory level over time. Figure 9.4 shows a simplified inventory profile for one particular stock item in a retail operation. Every time an order is placed, Q items are ordered. The replenishment order arrives in one batch instantaneously. Demand for the item is then steady and perfectly predictable at a rate of D units per

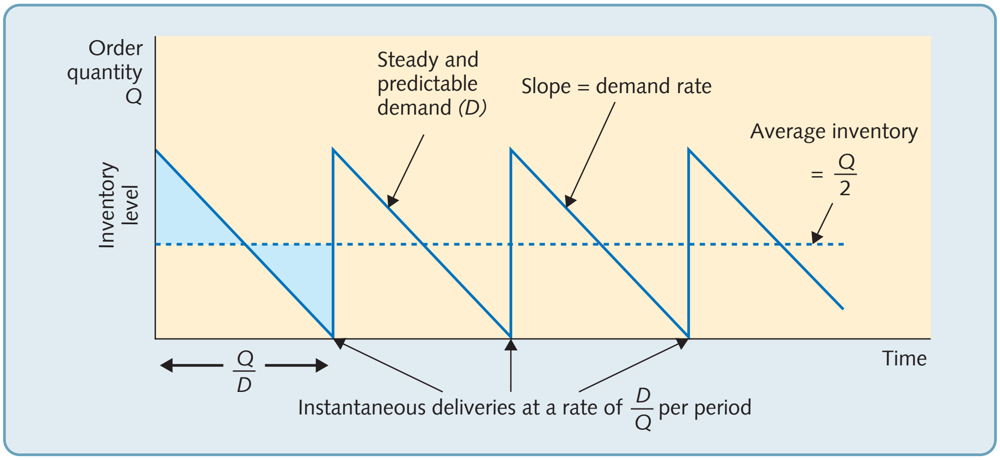

Figure 9.4 Inventory profiles chart the variation in inventory level

month. When demand has depleted the stock of the items entirely, another order of Q items instantaneously arrives, and so on. Under these circumstances:

The average inventory $= \frac{Q}{2}$ (because the two shaded areas in Figure 9.3 are equal)

The time interval between deliveries $= \frac{Q}{D}$

The frequency of deliveries = the reciprocal of the time interval $= \frac{D}{Q}$

The economic order quantity (EOQ) formula

The EOQ approach attempts to find the best balance between the advantages and disadvantages of holding stock. For example, Figure 9.5 shows two alternative order-quantity policies for an item. Plan A, represented by the unbroken line, involves ordering in quantities of 400 at a time. Demand in this case is running at 1,000 units per year. Plan B, represented by the dotted line, uses smaller but more frequent replenishment orders. This time only 100 are ordered at a time, with orders being placed four times as often. However, the average inventory for plan B is one-quarter of that for plan A.

To find out whether either of these plans, or some other plan, minimises the total cost of stocking the item, we need some further information, namely the total cost of holding one unit in stock for a period of time (C_h) and the total costs of placing an order (C_o).

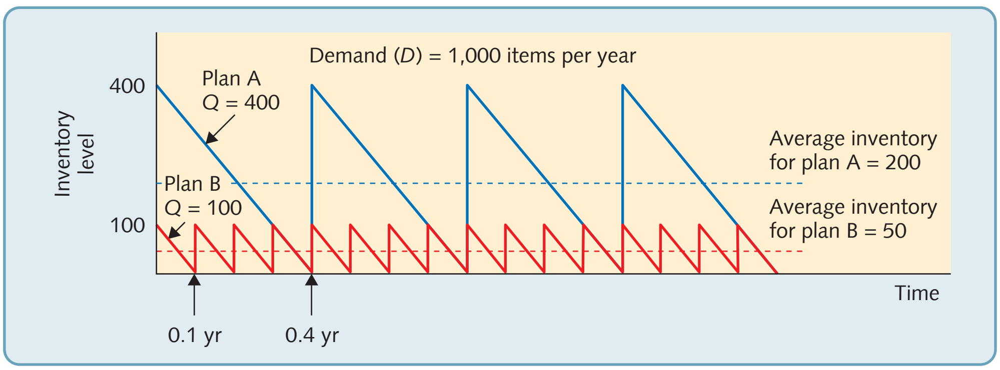

Figure 9.5 Two alternative inventory plans with different order quantities (*Q*)

In this case, the cost of holding stocks is calculated at £1 per item per year and the cost of placing an order is calculated at £20 per order.

We can now calculate total holding costs and ordering costs for any particular ordering plan as follows:

$$\text{Holding costs} = \text{holding cost/unit} \times \text{average inventory}$$
$$= C_h \times \frac{Q}{2}$$
$$\text{Ordering costs} = \text{ordering cost} \times \text{number of orders per period}$$
$$= C_o \times \frac{Q}{2}$$
$$\text{Total cost, } C_t = \frac{C_h Q}{2} + \frac{C_o D}{Q}$$

We can now calculate the costs of adopting plans with different order quantities. These are illustrated in Table 9.4. As we would expect with low values of Q, holding costs are low but ordering costs are high, because orders have to be placed very frequently. As Q increases, the holding costs increase but the costs of placing orders decrease. In this case, the order quantity, Q, which minimises the sum of holding and order costs, is 200. This 'optimum' order quantity is called the EOQ. This is illustrated graphically in Figure 9.6.

A more elegant method of finding the EOQ is to derive its general expression. This can be done using simple differential calculus as follows. From before:

$$\text{Total cost} = \text{holding cost} + \text{order cost}$$
$$C_t = \frac{C_h Q}{2} + \frac{C_o D}{Q}$$

The rate of change of total cost is given by the first differential of C_t with respect to Q:

$$\frac{dC_t}{dQ} = \frac{C_h}{2} - \frac{C_o D}{Q^2}$$

Table 9.4 **Costs of adoption of plans with different order quantities**

Demand (D) = 1,000 units per year   Holding costs (C_h) = £1 per item per year
Order costs (C_o) = £20 per order

Order quantity (*Q*)	*Holding costs* ($0.5Q \times C_h$)	+	*Order costs* ($(D/Q) \times C_o$)	=	*Total costs*
50	25		20 × 20 = 400		425
100	50		10 × 20 = 200		250
150	75		6.7 × 20 = 134		209
200	100		5 × 20 = 100		200*
250	125		4 × 20 = 80		205
300	150		3.3 × 20 = 66		216
350	175		2.9 × 20 = 58		233
400	200		2.5 × 20 = 50		250

*Minimum total cost.

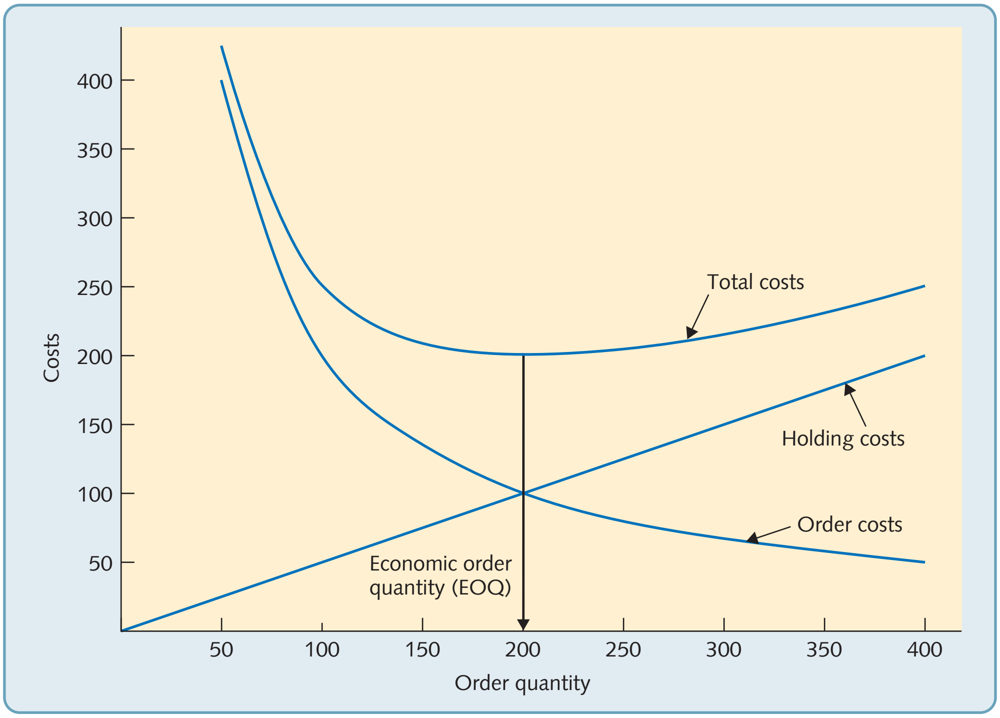

Figure 9.6 **Inventory-related costs minimise at the 'economic order quantity' (EOQ)**

The lowest cost will occur when $\frac{dC_t}{dQ} = 0$ that is:

$$0 = \frac{C_h}{2} - \frac{C_o D}{Q_o^2}$$

where Q_o = the EOQ. Rearranging this expression gives:

$$Q_o = \text{EOQ} = \sqrt{\frac{2C_o D}{C_h}}$$

When using the EOQ:

$$\text{Time between orders} = \frac{\text{EOQ}}{D}$$

$$\text{Order frequency} = \frac{D}{\text{EOQ}} \text{ per period}$$

Sensitivity of the EOQ

The graphical representation of the total cost curve in Figure 9.6 shows that, although there is a single value of Q which minimises total costs, any relatively small deviation from the EOQ will not increase total costs significantly. In other words, costs will be near-optimum provided a value of Q which is reasonably close to the EOQ is chosen. Put another way, small errors in estimating either holding costs or order costs will not result in a significant change in the EOQ. This is a particularly convenient phenomenon because, in practice, both holding and order costs are not easy to estimate accurately. The other implication is that, because the total cost curve is not symmetrical, it is usually better to have slightly more than slightly less inventory.

OPERATIONS PRINCIPLE

For any stock replenishment activity there is a theoretical 'optimum' order quantity that minimises total inventory-related costs.

WORKED EXAMPLE

A building materials supplier obtains its bagged cement from a single supplier. Demand is reasonably constant throughout the year, and last year the company sold 2,000 tonnes of this product. It estimates the costs of placing an order at around £25 each time an order is placed, and calculates that the annual cost of holding inventory is 20 per cent of purchase cost. The company purchases the cement at £60 per tonne. How much should the company order at a time?

$$\text{EOQ for cement} = \sqrt{\frac{2C_oD}{C_h}}$$

$$= \sqrt{\frac{2 \times 25 \times 2{,}000}{0.2 \times 60}}$$

$$= \sqrt{\frac{100{,}000}{12}}$$

$$= 91.287 \text{ tonnes}$$

After calculating the EOQ the operations manager feels that placing an order for 91.287 tonnes exactly seems somewhat over-precise. Why not order a convenient 100 tonnes?

Total cost of ordering plan for $Q = 91.287$:

$$= \frac{C_hQ}{2} + \frac{C_oD}{Q}$$

$$= \frac{(0.2 \times 60) \times 91.287}{2} + \frac{25 \times 2{,}000}{91.287}$$

$$= £1095.454$$

Total cost of ordering plan for $Q = 100$:

$$= \frac{(0.2 \times 60) \times 100}{2} + \frac{25 \times 2{,}000}{100}$$

$$= £1{,}100$$

The extra cost of ordering 100 tonnes at a time is £1,100 − £1,095.45 = £4.55. The operations manager therefore should feel confident in using the more convenient order quantity.

Gradual replacement – the economic batch quantity (EBQ) model

The simple inventory profile shown in Figure 9.4 assumes that each complete replacement order arrives at one point in time. However, replenishment may occur over a time period rather than in one lot, for example where an internal order is placed for a batch of parts to be produced on a machine. The machine will start to produce items and ship them in a more-or- less, continuous stream into inventory, but at the same time demand is removing items from the inventory. Provided the rate at which items are supplied to the inventory (P) is higher than the demand rate (D), then the inventory will increase. After the batch has been completed the machine will be reset (to produce some other part), and demand will continue to deplete the inventory level until production of the next batch begins. The resulting profile is shown in Figure 9.7. This is typical for inventories supplied by batch processes, and the minimum-cost batch quantity for this profile is called the economic batch quantity (EBQ). It is derived as follows:

$$\text{Maximum stock level} = M$$

$$\text{Slope of inventory build-up} = P - D$$

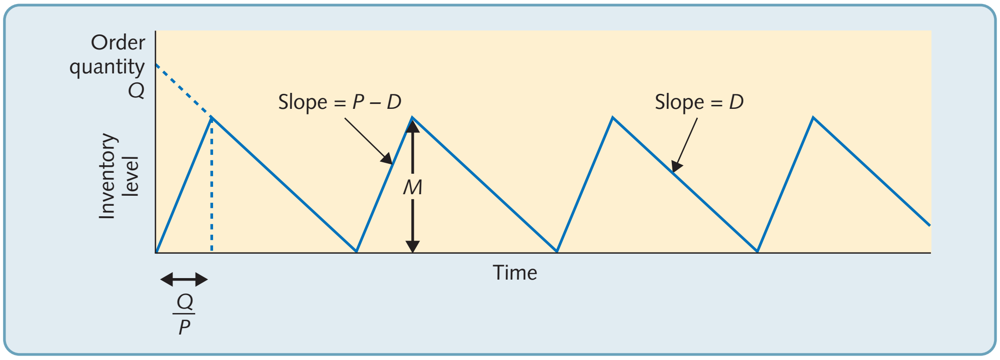

Figure 9.7 Inventory profile for gradual replacement of inventory

Also, as is clear from Figure 9.7:

$$\text{Slope of inventory build-up} = M \div \frac{Q}{P} = \frac{MP}{Q}$$

So,

$$\frac{MP}{Q} = P - D$$

$$M = \frac{Q(P - D)}{P}$$

$$\text{Average inventory level} = \frac{M}{2} = \frac{Q(P - D)}{2P}$$

As before:

$$\text{Total cost} = \text{holding cost} + \text{order cost}$$

$$C_t = \frac{C_h Q(P - D)}{2P} + \frac{C_o D}{Q}$$

$$\frac{dC_t}{dQ} = \frac{C_h(P - D)}{2P} - \frac{C_o D}{Q^2}$$

Again, equating to zero and solving Q gives the minimum-cost order quantity EBQ:

$$\text{EBQ} = \sqrt{\frac{2C_o D}{C_h(1 - (D/P))}}$$

WORKED EXAMPLE

The manager of a bottle-filling plant, which bottles soft drinks, needs to decide how long a 'run' of each type of drink to process. Demand for each type of drink is reasonably constant at 80,000 per month (a month has 160 production hours). The bottling lines fill at a rate of 3,000 bottles per hour, but take an hour to clean and reset between different drinks. The cost (of labour and lost production capacity) of each of these changeovers has been calculated at £100 per hour. Stock-holding costs are counted at £0.1 per bottle per month.

$$D = 80{,}000 \text{ per month} = 500 \text{ per hour}$$

$$\text{EBQ} = \sqrt{\frac{2C_o D}{C_h(1 - (D/P))}}$$

$$= \sqrt{\frac{2 \times 100 \times 80{,}000}{0.1(1 - (500/3{,}000))}}$$

$$\text{EBQ} = 13{,}856$$

The staff who operate the lines have devised a method of reducing the changeover time from 1 hour to 30 minutes. How would that change the EBQ?

$$\text{New Co} = £50$$

$$\text{New EBQ} = \sqrt{\frac{2 \times 50 \times 80{,}000}{0.1(1 - (500/3{,}000))}}$$

$$= 9{,}798$$

If customers won't wait – the newsvendor problem

A special case of the inventory order-quantity decision is when an order quantity is purchased for a specific event or time period, after which the items are unlikely to be sold. A simple example of this is the decision taken by a newspaper vendor of how many newspapers to stock for the day. If the newsvendor should run out of papers, customers will either go elsewhere or decide not to buy a paper that day. Newspapers left over at the end of the day are worthless and demand for the newspapers varies day-by-day. In deciding how many newspapers to carry, the newsvendor is, in effect, balancing the risk and consequence of running out of newspapers against that of having newspapers left over at the end of the day. Retailers and manufacturers of high-class leisure products, such as some books and popular music CDs, face the same problem. For example, a concert promoter needs to decide how many concert T-shirts to order emblazoned with the logo of the main act. The profit on each T-shirt sold at the concert is £5 and any unsold T-shirts are returned to the company that supplies them, but at a loss to the promoter of £3 per T-shirt. Demand is uncertain but is estimated to be between 200 and 1,000. The probabilities of different demand are as follows:

Demand level	200	400	600	800
Probability	0.2	0.3	0.4	0.1

How many T-shirts should the promoter order? Table 9.5 shows the profit that the promoter would make for different order quantities and different levels of demand.

We can now calculate the expected profit that the promoter will make for each order quantity by weighting the outcomes by their probability of occurring.

If the promoter orders 200 T-shirts:

$$\text{Expected profit} = 1{,}000 \times 0.2 + 1{,}000 \times 0.3 + 1{,}000 \times 0.4 + 1{,}000 \times 0.1$$
$$= £1{,}000$$

Table 9.5 Pay-off matrix for T-shirt order quantity (profit or loss in £s)

Demand level	*200*	*400*	*600*	*800*
Probability	0.2	0.3	0.4	0.1
Promoter orders 200	1,000	1,000	1,000	1,000
Promoter orders 400	400	2,000	2,000	2,000
Promoter orders 600	−200	1,400	3,000	3,000
Promoter orders 800	−800	800	2,400	4,000

If the promoter orders 400 T-shirts:

$$\text{Expected profit} = 400 \times 0.2 + 2{,}000 \times 0.3 + 2{,}000 \times 0.4 + 2{,}000 \times 0.1 = £1{,}680$$

If the promoter orders 600 T-shirts:

$$\text{Expected profit} = -200 \times 0.2 + 1{,}400 \times 0.3 + 3{,}000 \times 0.4 + 3{,}000 \times 0.1 = £1{,}880$$

If the promoter orders 800 T-shirts:

$$\text{Expected profit} = -800 \times 0.2 + 800 \times 0.3 + 2{,}400 \times 0.4 + 4{,}000 \times 0.1 = £1{,}440$$

The order quantity that gives the maximum profit is 600 T-shirts, which results in a profit of £1880.

The importance of this approach lies in the way it takes a probabilistic view of part of the inventory calculation (demand). Something we shall use again in this chapter

EXAMPLE

Mr Ruben's bakery[4]

Be careful about treating the newsvendor problem on a product-by-product basis. It is a powerful idea, but needs to be seen in context. Take the famous (in New York) City Bakery, in Manhattan, which is run by master baker Maury Rubin, who knows the economics of baking fresh products. Ingredients and rent are expensive. It costs Mr Rubin $2.60 to make a $3.50 croissant. If he makes 100 and sells 70, he earns $245 but his costs are $260, and because all goods are sold within a day (his quality standards mean that he will not sell leftovers), he loses money. Nor can he raise his prices. In his competitive market, he says, shoppers bristle when the cost of baked goods passes a certain threshold. However, Mr Ruben has two 'solutions'. First, he can subsidise his croissants by selling higher-margin items such as fancy salads and sandwiches. Second, he uses data to cut waste, by studying sales so that he can detect demand trends in order to fine-tune supply. He monitors the weather carefully (demand drops away when it rains) and carefully inspects school calendars so he can reduce the quantities he bakes during school holidays. Each day in the morning, he makes sure that pastries are prepared, but then he checks sales every 60–90 minutes before making the decision to adjust supply or not. Only when the numbers are in do the pastries go into the oven. Having no croissants left by the end of the day is a sign of success.

DIAGNOSTIC QUESTION

Are inventory orders being placed at the right time?

When we assumed that orders arrived instantaneously and demand was steady and predictable, the decision on when to place a replenishment order was self-evident. An order would be placed as soon as the stock level reached zero, it would arrive instantaneously and prevent any stock-out occurring. When there is a lag between the order being placed and it arriving in the inventory, we can still calculate the timing of a replacement order simply, as shown in Figure 9.8. The lead time for an order to arrive is in this case two weeks, so the reorder point (ROP) is the point at which stock will fall to zero, minus the order lead time. Alternatively, we can define the point in terms of the level that the inventory will have reached when a replenishment order needs to be placed. Here, this occurs at a reorder level (ROL) of 200 items.

However, this assumes that both the demand and the order lead time are perfectly predictable. In most cases this is not so. Both demand and the order lead time are likely to vary to produce a profile that looks something like that in Figure 9.9. In these circumstances, it is necessary to make the replenishment order somewhat earlier than would be the case in a purely deterministic situation. This will result in, on average, some 'safety' stock still being in the inventory when the replenishment order arrives. The earlier the replenishment order is placed, the higher will be the expected level of safety stock when the replenishment order arrives. However, because of the variability of both lead time (t) and demand rate (d), the safety stock at the time of replenishment will vary. The main consideration in setting safety stock is the probability that the stock will not have run out before the replenishment order arrives. This depends on the lead-time usage distribution. This is a combination of the distributions that describe lead-time variation and the demand rate during the lead time. If safety stock is set below the lower limit of this distribution then there will be shortages every single replenishment cycle. If safety stock is set above the upper limit of the distribution, there is no chance of

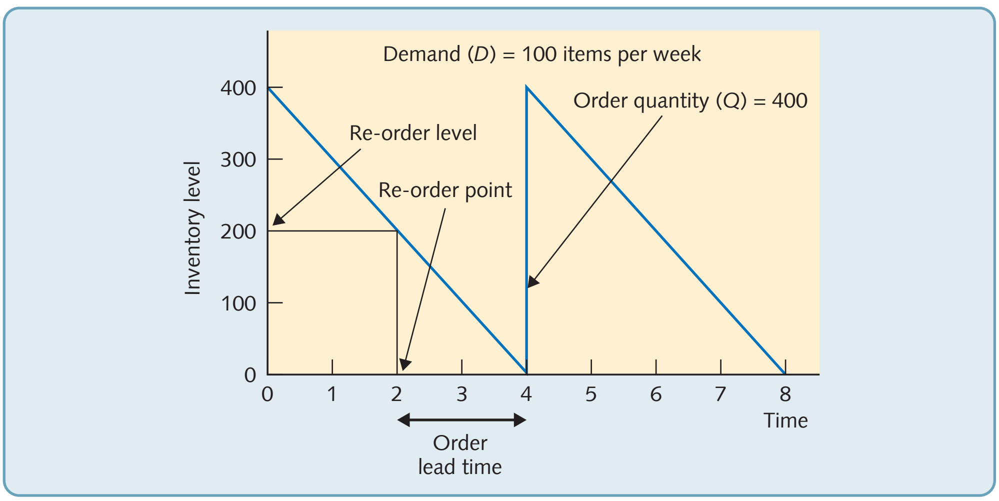

Figure 9.8 **Reorder level (ROL) and reorder point (ROP) are derived from the order lead time and demand rate**

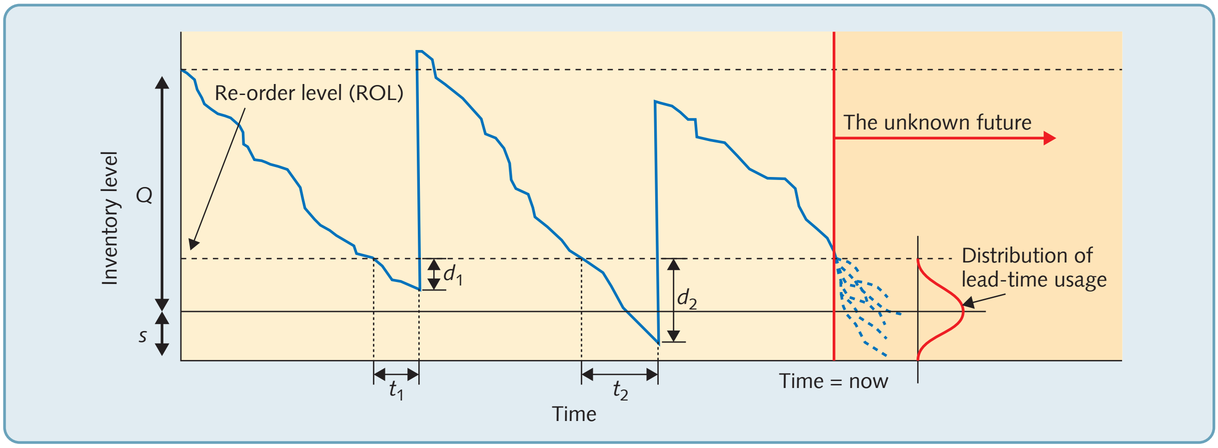

Figure 9.9 **Safety stock (*s*) helps to avoid stock-outs when demand and/or order lead time are uncertain**

> **OPERATIONS PRINCIPLE**
>
> ***For any stock replenishment activity, the timing of replenishment should reflect the effects of uncertain lead time and uncertain demand during that lead time.***

stock-outs occurring. Usually, safety stock is set to give a predetermined likelihood that stock-outs will not occur. Figure 9.9 shows that, in this case, the first replenishment order arrived after t_1, resulting in a lead-time usage of d_1. The second replenishment order took longer, t_2, and demand rate was also higher, resulting in a lead-time usage of d_2. The third-order cycle shows several possible inventory profiles for different conditions of lead-time usage and demand rate.

WORKED EXAMPLE

An online retailer of running shoes can never be certain of how long, after placing an order, the delivery will take. Examination of previous orders reveals that out of 10 orders: one took 1 week, two took 2 weeks, four took 3 weeks, two took 4 weeks and one took 5 weeks. The rate of demand for the shoes also varies between 110 pairs per week and 140 pairs per week. There is a 0.2 probability of the demand rate being either 110 or 140 pairs per week, and a 0.3 chance of demand being either 120 or 130 pairs per week. The company needs to decide when it should place replenishment orders if the probability of a stock-out is to be less than 10 per cent.

Both lead time and the demand rate during the lead time will contribute to the lead-time usage. So the distributions that describe each will need to be combined. Figure 9.10 and Table 9.5 show how this can be done. Taking lead time to be 1, 2, 3, 4 or 5 weeks, and demand rate to be 110, 120, 130 or 140 pairs per week, and also assuming the two variables to be independent, the distributions can be combined as shown in Table 9.6. Each element in the matrix shows a possible lead-time usage with the probability of its occurrence. So if the lead time is one week and the demand rate is 110 pairs per week, the actual lead-time usage will be $1 \times 110 = 110$ pairs. Since there is a 0.1 chance of the lead time being one week, and a 0.2 chance of demand rate being 110 pairs per week, the probability of both these events occurring is $0.1 \times 0.2 = 0.02$.

We can now classify the possible lead-time usages into histogram form. For example, summing the probabilities of all the lead-time usages which fall within the range 100–199 (all the first column) gives a combined probability of 0.1. Repeating this for subsequent intervals results in Table 9.7.

This shows the probability of each possible range of lead-time usage occurring, but it is the cumulative probabilities that are needed to predict the likelihood of stock-out (see Table 9.8).

Setting the reorder level at 600 would mean that there is only a 0.08 chance of usage being greater than available inventory during the lead time, that is, there is a less than 10 per cent chance of a stock-out occurring.

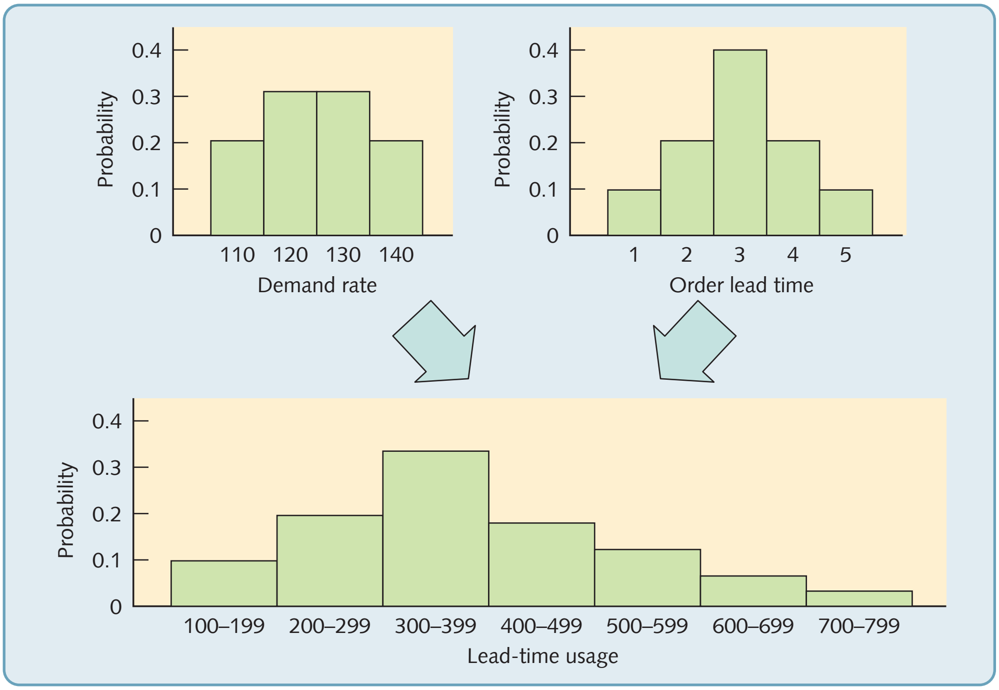

Figure 9.10 **The probability distributions for order lead time and demand rate combine to give the lead-time usage distribution**

Table 9.6 Matrix of lead time and demand-rate probabilities

			Lead-time probabilities				
			1	*2*	*3*	*4*	*5*
			0.1	*0.2*	*0.4*	*0.2*	*0.1*
Demand-rate probabilities	110	0.2	110 (0.02)	220 (0.04)	330 (0.08)	440 (0.04)	550 (0.02)
	120	0.3	120 (0.03)	240 (0.06)	360 (0.12)	480 (0.06)	600 (0.03)
	130	0.3	130 (0.03)	260 (0.06)	390 (0.12)	520 (0.06)	650 (0.03)
	140	0.2	140 (0.02)	280 (0.04)	420 (0.08)	560 (0.04)	700 (0.02)

Table 9.7 Combined probabilities

Lead-time usage	*100–199*	*200–299*	*300–399*	*400–499*	*500–599*	*600–699*	*700–799*
Probability	0.1	0.2	0.32	0.18	0.12	0.06	0.02

Table 9.8 Combined probabilities

Lead-time usage	*100*	*200*	*300*	*400*	*500*	*600*	*700*	*800*
Probability of usage being greater than *X*	01.0	0.9	0.7	0.38	0.2	0.08	0.02	0

Continuous and periodic review

The approach we have described is often called the continuous review approach. To make the decision in this way the stock level of each item must be reviewed continuously and an order placed when the stock level reaches its reorder level. The virtue of this approach is that, although the timing of orders may be irregular (depending on the variation in demand rate), the order size (Q) is constant and can be set at the optimum economic order quantity. But continually checking inventory levels may be time-consuming. An alternative, and simpler, approach, but one that sacrifices the use of a fixed (and therefore possibly optimum) order quantity, is 'periodic review'. Here, rather than ordering at a predetermined reorder level, the periodic approach orders at a fixed and regular time interval. So the stock level of an item could be found, for example, at the end of every month and a replenishment order placed to bring the stock up to a predetermined level. This level is calculated to cover demand between the replenishment order being placed and the following replenishment order arriving. Plus safety stocks will need to be calculated, in a similar manner to before, based on the distribution of usage over this period.

Two-bin and three-bin systems

Keeping track of inventory levels is especially important in continuous review approaches to reordering. A simple and obvious method of indicating when the reorder point has been reached is necessary, especially if there are a large number of items to be monitored. The simple two-bin system involves storing the reorder point quantity plus the safety inventory quantity in the second bin and using parts from the first bin. When the first bin empties, it is the signal to order the next reorder quantity. Different 'bins' are not always necessary to operate this type of system. For example, a common practice in retail operations is to store the second 'bin' quantity upside-down behind or under the first 'bin' quantity. Orders are then placed when the upside-down items are reached.

DIAGNOSTIC QUESTION

Is inventory being controlled effectively?

Even probabilistic models are still simplified compared with the complexity of real stock management. Coping with many thousands of stocked items, supplied by many hundreds of different suppliers, with possibly tens of thousands of individual customers, makes for a complex and dynamic operations task. Controlling such complexity requires an approach that discriminates between different items so that each has a degree of control that is appropriate to its importance. It also requires an information system to keep track of inventories.

Inventory priorities – the ABC system

Some stocked items are more important than others. Some might have a high usage rate, so if they ran out many customers would be disappointed. Others might be of particularly high value, so excessively high inventory levels would be particularly expensive. One common way of discriminating between different stock items is to rank them by their usage value (usage rate multiplied by value). Items with a particularly high-usage value are deemed to warrant the most careful control, whereas those with low-usage values need not be controlled quite so rigorously. Generally, a relatively small proportion of the total range of items contained in an inventory will account for a large proportion of the total usage value. This phenomenon is known as the Pareto, or 80/20 rule. It is so called because, typically, 80 per cent of an operation's sales are accounted for by only 20 per cent of all stocked item types. (This idea is also used elsewhere in operations management; for example, see Chapter 12.) Here the relationship is used to classify items into A, B, or C categories, depending on their usage value:

OPERATIONS PRINCIPLE

Different inventory management decision rules are needed for different classes of inventory.

- Class A items are those 20 per cent or so of high-usage value items that account for around 80 per cent of the total usage value.
- Class B items are those of medium-usage value, usually the next 30 per cent of items, which often account for around 10 per cent of the total usage value.
- Class C items are those low-usage value items which, although comprising around 50 per cent of the total types of items stocked, probably only account for around 10 per cent of the total usage value of the operation.

Although annual usage and value are the two criteria most commonly used to determine a stock classification system, other criteria might also contribute towards the (higher) classification of an item. The consequence of stock-out might give higher priority to some items that would seriously delay or disrupt operations, if they were not in stock. Uncertainty of supply may also give some items priority, as might high obsolescence or deterioration risk.

WORKED EXAMPLE

Table 9.9 shows all the parts stored by an electrical wholesaler. The 20 different items stored vary in terms of both their usage per year and cost per item as shown. However, the wholesaler has ranked the stock items by their usage value per year. The total usage value per year is

Table 9.9 Warehouse items ranked by usage value

Stock no.	*Usage (items/year)*	*Cost (£/item)*	*Usage value (£000/year)*	*% of total value*	*Cumulative % of total value*
A/703	700	20.00	14,000	25.41	25.14
D/012	450	2.75	1238	22.23	47.37
A/135	1,000	0.90	900	16.16	63.53
C/732	95	8.50	808	14.51	78.04
C/735	520	0.54	281	5.05	83.09
A/500	73	2.30	168	3.02	86.11
D/111	520	0.22	114	2.05	88.16
D/231	170	0.65	111	1.99	90.15
E/781	250	0.34	85	1.53	91.68
A/138	250	0.30	75	1.34	93.02
D/175	400	0.14	56	1.01	94.03
E/001	80	0.63	50	0.89	94.92
C/150	230	0.21	48	0.86	95.78

(*contined*)

Table 9.9 Warehouse items ranked by usage value (contined)

Stock no.	*Usage (items/year)*	*Cost (£/item)*	*Usage value (£000/year)*	*% of total value*	*Cumulative % of total value*
F/030	400	0.12	48	0.86	96.64
D/703	500	0.09	45	0.81	97.45
D/535	50	0.88	44	0.79	98.24
C/541	70	0.57	40	0.71	98.95
A/260	50	0.64	32	0.57	99.52
B/141	50	0.32	16	0.28	99.80
D/021	20	0.50	10	0.20	100.00
Total			5,569	100.00	

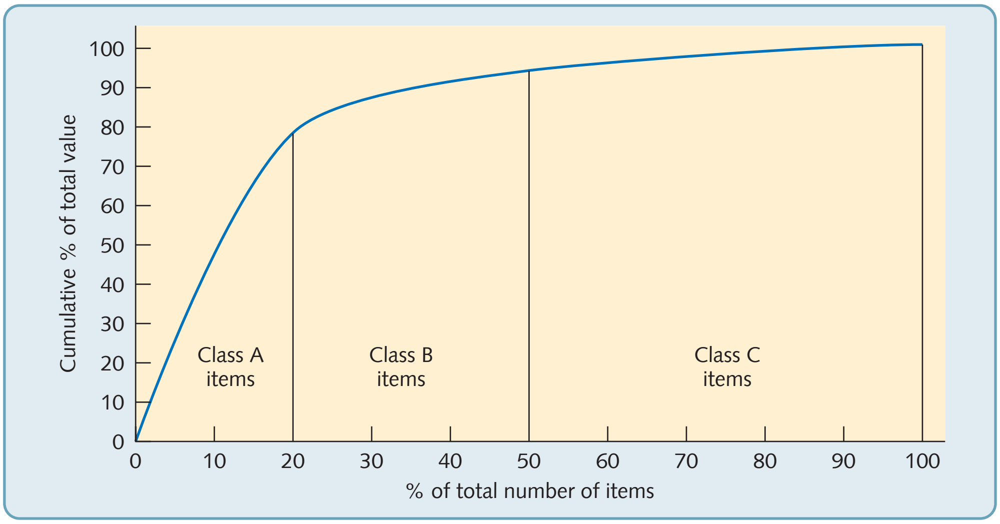

Figure 9.11 Pareto curve for items in a warehouse

£5,569,000. From this, it is possible to calculate the usage value per year of each item as a percentage of the total usage value, and from that a running cumulative total of the usage value as shown. The wholesaler can then plot the cumulative percentage of all stocked items against the cumulative percentage of their value. For example, the part with stock number A/703 is the highest value part and accounts for 25.14 per cent of the total inventory value. As a part, however, it is only one-twentieth or 5 per cent of the total number of items stocked. This item together with the next highest value item (D/012) account for only 10 per cent of the total number of items stocked, yet account for 47.37 per cent of the value of the stock, and so on.

These data are shown graphically in Figure 9.11. The first four part numbers (20 per cent of the range) are Class A items, whose usage will be monitored very closely. The six next part numbers (30 per cent of the range) are Class B items with slightly less effort devoted to their control. All other items are Class C items, whose stocking policy is reviewed only occasionally.

Inventory information systems

Most inventories of any significant size are managed by information systems. This is especially so since the use of bar-code readers, radio frequency identification (RFID) and the point-of-sale recording of sales transactions has made data capture more convenient. Many commercial systems of stock control are available, although they tend to share certain common functions.

Updating stock records

Every time an inventory transaction takes place the position, status and possibly value of the stock will have changed. This information must be recorded so that operations managers can determine their current inventory status at any time.

Generating orders

Both the how much and the when to order decisions, can be made by a stock control system. Originally, almost all computer systems calculated order quantities by using the EOQ formulae. Now more sophisticated probabilistic algorithms are used, based on examining the marginal return on investing in stock. The system will hold all the information which goes into the ordering algorithm but might periodically check to see if demand or order lead times, or any of the other parameters, have changed significantly and recalculate accordingly. The decision on when to order, on the other hand, is a far more routine affair, which computer systems make according to whatever decision rules operations managers have chosen to adopt: either continuous review or periodic review.

Generating inventory reports

Inventory control systems can generate regular reports of stock value, which can help management monitor its inventory control performance. Similarly, customer service performance, such as the number of stock-outs or the number of incomplete orders, can be regularly monitored. Some reports may be generated on an exception basis. That is, the report is only generated if some performance measure deviates from acceptable limits.

Forecasting

Ideally, inventory replenishment decisions should be made with a clear understanding of forecast future demand. Inventory control systems usually compare actual demand against forecast and adjust forecasts in the light of actual levels of demand.

Common problems with inventory systems

Our description of inventory systems has been based on the assumption that operations have:

- a reasonably accurate idea of costs such as holding cost, or order cost, and
- accurate information that really does indicate the actual level of stock and sales.

In fact, data inaccuracy often poses one of the most significant problems for inventory managers. This is because most computer-based inventory management systems are based on what is called the perpetual inventory principle. This is the simple idea that stock records are (or should be) automatically updated every time items are recorded as having been received into an inventory or taken out of the inventory. So:

$$\text{opening stock level} + \text{receipts in} - \text{dispatches out} = \text{new stock level}$$

OPERATIONS PRINCIPLE

The maintenance of data accuracy is vital for the day-to-day effectiveness of inventory management systems.

Any errors in recording these transactions, and/or in handling the physical inventory can lead to discrepancies between the recorded and actual inventory, and these errors are perpetuated until physical stock checks are made (usually quite infrequently). In practice, there are many opportunities for errors to occur, if only because inventory transactions are numerous. This means that it is surprisingly common for the majority of inventory records to be inaccurate. The underlying causes of errors include:

- keying errors; entering the wrong product code
- quantity errors; a miscount of items put into or taken from stock
- damaged or deteriorated inventory not recorded as such, or not correctly deleted from the records when it is destroyed
- the wrong items being taken out of stock, but the records not being corrected when they are returned to stock
- delays between the transactions being made and the records being updated
- items stolen from inventory (common in retail environments, but also not unusual in industrial and commercial inventories).

Critical commentary

The approach to determining order quantity that involves optimising costs of holding stock against costs of ordering stock, typified by the EOQ and EBQ models, has always been subject to criticisms. Originally, these concerned the validity of some of the assumptions of the model; more recently, they have involved the underlying rationale of the approach itself. Criticisms include: the assumptions included in the EOQ models are simplistic; the real costs of stock in operations are not as assumed in EOQ models; and cost minimisation is not an appropriate objective for inventory management.

- The last criticism is particularly significant. Many organisations (such as supermarkets and wholesalers) make the most of their revenue and profits simply by holding and supplying inventory. Because their main investments are in the inventory, it is critical that they make a good return on this capital, by ensuring that it has the highest possible 'stock turn' and/or gross profit margin. Alternatively, they may also be concerned to maximise the use of space by seeking to maximise the profit earned per square metre. The EOQ model does not address these objectives. Similarly, for products that deteriorate or go out of fashion, the EOQ model can result in excess inventory of slower-moving items. In fact, the EOQ model is rarely used in such organisations, and there is more likely to be a system of periodic review for regular ordering of replenishment inventory. For example, a typical builders' supply merchant might carry around 50,000 different items of stock (SKUs). However, most of these cluster into larger families of items such as paints, sanitary ware, or metal fixings. Single orders are placed at regular intervals for all the required replenishments in the supplier's range, and these are then delivered together at one time. If deliveries are made weekly then, on average, the individual item order quantities will be for only one-week's usage. Less popular items, or ones with erratic demand patterns, can be ordered individually at the same time, or (when urgent) can be delivered the next day by carrier.

- Some regard the ABC approach to inventory classification as misleading. Many professional inventory managers point out that the slow-moving (C category) items often pose the greatest challenge in inventory management. Often these slow-moving items, although only accounting for 20 per cent of sales, require a large part (typically between one-half and two-thirds) of the total investment in stock. This is why slow-moving items are a real problem. Moreover, if errors in forecasting or ordering result in excess stock in 'A class' fast-moving items, it is relatively unimportant in the sense that excess stock can be sold quickly. However, excess stock in slow-moving C items will be there for a long time. According to some inventory managers, it is the A items that can be left to look after themselves, it is the B and even more the C items that need controlling.

SUMMARY CHECKLIST

- ☐ Have all inventories been itemised and costed?
- ☐ Have all the costs and negative effects of inventory been assessed?
- ☐ What proportion of inventory is there:
  - – as an insurance against uncertainty?
  - – to counteract a lack of flexibility?
  - – to allow operations to take advantage of short-term opportunities?
  - – to anticipate future demand?
  - – to reduce overall costs?
  - – because it can increase in value?
  - – because it is in the processing pipeline?
- ☐ Have methods of reducing inventory in these categories been explored?
- ☐ Have cost minimisation methods been used to determine order quantity?
- ☐ Do these use a probabilistic estimate of demand?
- ☐ Have the relative merits of continuous and period inventory review been assessed?
- ☐ Are probabilistic estimates of demand and lead time used to determine safety stock levels?
- ☐ Are items controlled by their usage value?
- ☐ Does the inventory information system integrate all inventory decisions?

CASE STUDY

supplies4medics.com

Founded at the height of the 'dotcom bubble' of the late 1990s, supplies4medics.com has become one of Europe's most successful direct mail suppliers of medical hardware and consumables to hospitals, doctors' and dentists' surgeries, clinics, nursing homes and other medical-related organisations. Its physical and online catalogues list just over 4,000 items, categorised by broad applications such as 'hygiene consumables' and 'surgeons' instruments'. Quoting their web site:

> *'We are the pan-European distributors of wholesale medical and safety supplies. . . We aim to carry everything you might ever need: from nurses' scrubs to medical kits, consumables for operations, first aid kits, safety products, chemicals, fire-fighting equipment, nurse and physicians' supplies, etc. Everything is at affordable prices – and backed by our very superior customer service and support – supplies4medics is your ideal source for all medical supplies. Orders are normally dispatched same-day, via our European distribution partner, the Brussels Hub of DHL. You should therefore receive your complete order within one week, but you can request next day delivery if required, for a small extra charge. You can order our printed catalogue on the link at the bottom of this page, or shop on our easy-to-use online store.'*

Last year turnover grew by over 25 per cent to about €120 million, a cause for considerable satisfaction in the company. However, profit growth was less spectacular; and market research suggested that customer satisfaction, although generally good, was slowly declining. Most worrying, inventory levels had grown faster than sales revenue, in percentage terms. This was putting a strain on cash flow, requiring the company to borrow more cash to fund the rapid growth planned for the next year. Inventory holding is estimated to be costing around 15 per cent per annum, taking account of the cost of borrowing, insurance, and all warehousing overheads.

Pierre Lamouche, the Head of Operations summarised the situation faced by his department:

> *'As a matter of urgency, we are reviewing our purchasing and inventory management systems! Most of our existing reorder levels (ROL) and reorder quantities (ROQ) were set several years ago, and have never been recalculated. Our focus has been on rapid growth through the introduction of new product lines. For more recently introduced items, the ROQs were based only on forecast sales, which actually can be quite misleading. We estimate that it costs us, on average, €50 to place and administer every purchase order, since most suppliers are still not able to take orders over the internet or by EDI. In the meantime, sales of some products have grown fast, while others have declined. Our average inventory (stock) cover is about 10 weeks, but . . . amazingly . . . we still run out of critical items! In fact, on average, we are currently out of stock of about 500 SKUs (Stock Keeping Units) at any time. As you can imagine, our service level is not always satisfactory with this situation. We really need help to conduct a review of our system, so have employed a mature intern from the local business school to review our system. He has first asked my team to provide information on a random, representative sample of 20 items from the full catalogue range, which is copied below' See Table 9.10.*

Table 9.10 Representative sample of 20 catalogue items

Sample number	*Catalogue reference number**	*Sales unit description***	*Sales unit cost(€)*	*Last 12 months' sales (units)*	*Inventory as at last year end (units)*	*Reorder quantity (units)*
1	11036	Disposable aprons (10pk)	2.40	100	0	10
2	11456	Ear-loop masks (Box)	3.60	6,000	1200	1,000
3	11563	Drill type 164	1.10	220	420	250
4	12054	Incontinence pads large	3.50	35,400	8,500	10,000
5	12372	150 ml syringe	11.30	430	120	100
6	12774	Rectal speculum 3 prong	17.40	65	20	20
7	12979	Pocket organiser blue	7.00	120	160	500
8	13063	Oxygen trauma kit	187.00	40	2	10
9	13236	Zinc oxide tape	1.50	1,260	0	50
10	13454	Dual head stethoscope	6.25	10	16	25
11	13597	Disp. latex catheter	0.60	3,560	12	20
12	13999	Roll-up wheelchair ramp	152.50	12	44	50
13	14068	WashClene tube	1.40	22,500	10,500	8,000
14	14242	Cervical collar	12.00	140	24	20
15	14310	Head wedge	89.00	44	2	10
16	14405	Three-wheel scooter	755.00	14	5	5
17	14456	Neonatal trach. tube	80.40	268	6	100
18	14675	Mouldable strip paste	10.20	1,250	172	100
19	14854	Sequential comp. pump	430.00	430	40	50
20	24943	Toilet safety frame	25.60	560	18	20

*Reference numbers are allocated sequentially as new items are added to catalogue.
**All quantities are in sales units (e.g. item, box, case, pack).

QUESTIONS

1 Prepare a spreadsheet-based ABC analysis of usage value. Classify as follows:

A Items: top 20 per cent of usage value
B Items: next 30 per cent of usage value
C Items: remaining 50 per cent of usage value

2 Calculate the inventory weeks for each item, for each classification, and for all the items in total. Does this suggest that the operations manager's estimate of inventory weeks is correct? If so, what is your estimate of the overall inventory at the end of the base year, and how much might that have increased during the year?

3 Based on the sample, analyse the underlying causes of the availability problem described in the text.

4 Calculate the EOQs for the A Items.

5 What recommendations would you give to the company?

APPLYING THE PRINCIPLES

Some of these exercises can be answered by reading the chapter. Others will require some general knowledge of business activity and some might require an element of investigation. **Model answers for the first two questions are available on the student companion website.** For model answers to the other questions in this section, please ask your tutor/lecturer.

1 A supplier makes monthly shipments to 'House & Garden Stores', in average lot sizes of 200 coffee tables. The average demand for these items is 50 tables per week, and the lead time from the supplier is 3 weeks. 'House & Garden Stores' must pay for inventory from the moment the supplier ships the products. If they are willing to increase their lot size to 300 units, the supplier will offer a lead time of 1 week. What will be the effect on cycle and pipeline inventories?

2 A local shop has a relatively stable demand for tins of sweetcorn throughout the year, with an annual total of 1,400 tins. The cost of placing an order is estimated at £15 and the annual cost of holding inventory is estimated at 25 per cent of the product's value. The company purchases tins for 20p. How much should the shop order at a time, and what is the total cost of the plan?

3 A fruit canning plant has a single line for three different fruit types. Demand for each type of tin is reasonably constant at 50,000 per month (a month has 160 production hours). The tinning process rate is 1,200 per hour, but it takes 2 hours to clean and reset between different runs. The cost of these changeovers (C_o) is calculated at £250 per hour. Stock-holding is calculated at £0.1 per tin per month. How big should the batch size be?

4 '*Our suppliers often offer better prices if we are willing to buy in larger quantities. This creates a pressure on us to hold higher levels of stock. Therefore, to find the best quantity to order we must compare the advantages of lower prices for purchases and fewer orders with the disadvantages of increased holding costs. This means that calculating total annual inventory-related costs should now not only include holding costs and ordering costs, but also the cost of purchased items themselves*.' (Manager, Tufton Bufton Port Importers Inc.). One supplier to Tufton Bufton Port Importers Inc. (TBPI) has introduced quantity discounts to encourage larger order quantities. The discounts are show below:

Order quantity	*Price per bottle (€)*
0–100	15.00
101–250	13.50
250^+	11.00

TBPI estimates that its annual demand for this particular wine is 1,500 bottles, its ordering costs are €30 per order, and its annual holding costs are 20 per cent of the bottle's price.

(a) How should TBPI go about deciding how many to order?

(b) How many should they order?

5 Revisit the example on the Blood and Transplant service at the beginning of the chapter.

(a) What are the factors that constitute inventory holding costs, order costs and stock-out costs in a National Blood Service?

(b) What makes this particular inventory planning and control example so complex?

(c) How might the efficiency with which a National Blood Service controls its inventory affect its ability to collect blood?

Notes on chapter

1 Sources include: *BBC News* website (2011) 'Blood bank "perfect storm" threat for 2012', 28 December; Stanger, S., Wilding, R., Yates, N., Cotton, S. (2012) 'What drives perishable inventory management performance? Lessons learnt from the UK blood supply chain', *Supply Chain Management*: An International Journal, vol. 17, Iss: 2, 107–123.

2 Sources include: *BBC News* (2012) 'How do they know when to grit roads?', 18 December; Grimm, E. (2009) 'Prepare for winter', *Daily Chronicle,* 30 November.

3 Sources include: Calder, S. (2017) 'British Airways faces £300,000 bill for flight to Barbados without enough toilet paper', *The Independent,* 7 March; Morris, H. (2017) 'British Airways faces £290,000 bill after running out of toilet roll', *The Telegraph,* 8 March.

4 *The Economist* (2015) 'Croissantonomics - Lessons in managing supply and demand for perishable products', *The Economist* Print Edition, 29 August.

TAKING IT FURTHER

Bragg, S. M. (2013) ***Inventory Management Paperback*****, Accounting Tools.** *A financial approach to the subject.*

Emmett, S. and Granville, D. (2007) ***Excellence in Inventory Management: How to Minimise Costs and Maximise Service, Cambridge Academic.*** *Practical and thorough examination of professional inventory management.*

Muller, M. (2011) ***Essentials of Inventory Management (2nd edn), Amacom*****.** *A good introduction.*

Relph, G. and Milner, C. (2015) ***Inventory Management: Advanced Methods for Managing Inventory within Business Systems, Kogan Page*****.** *Cuts through the complexity of inventory theory.*

10

Resource planning and control

Introduction

After the nature of an operation's resources have been, at least approximately, determined, the activities that create its services and products must be managed on an ongoing basis. It is concerned with what activities happen in operations, when they happen, where they happen and what resources are going to be involved. It ensures that materials or information or customers flow smoothly through processes, operations and supply networks, and that value-adding resources are managed efficiently and to avoid unnecessary delay. This is the activity of resource planning and control. It is a subject with many technical issues. We cover the best known of these (materials requirements planning (MRP) in the supplement to this chapter. Figure 10.1 shows the position of the ideas described in this chapter in the general model of operations management.

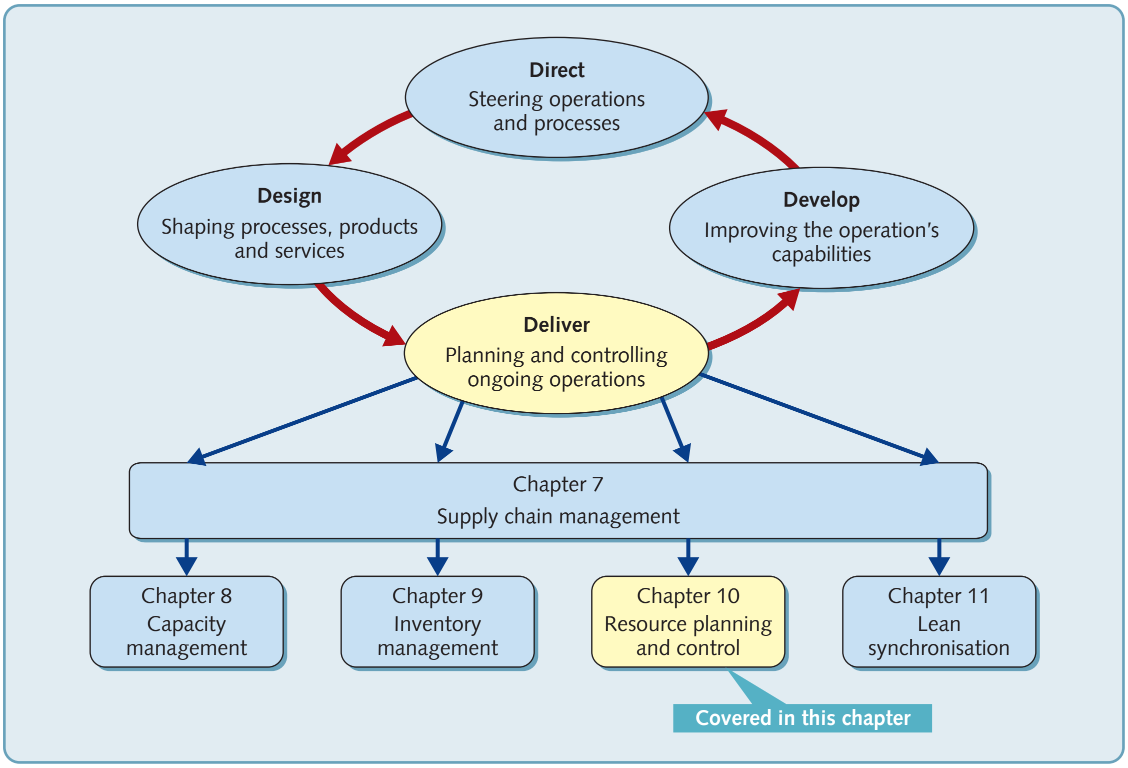

Figure 10.1 Resource planning and control is concerned with managing the ongoing allocation of resources and activities to ensure that the operation's processes are both efficient and reflect customer demand for products and services

EXECUTIVE SUMMARY

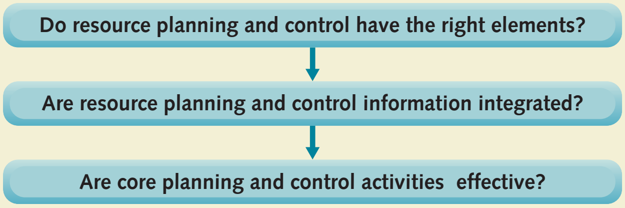

Do resource planning and control have all the right elements?

Resource planning and control is concerned with managing the ongoing allocation of resources and activities to ensure that the operation's processes are both efficient and reflect customer demand for products and services. In practice, planning (deciding what is intended to happen) and control (coping when things do not happen as intended), overlap to such an extent that they are usually treated together.

Although planning and control systems differ, they tend to have a number of common elements. These are: a customer interface that forms a two-way information link between the operation's activities and its customers; a supply interface that does the same thing for the operation's suppliers; a set of overlapping 'core' mechanisms that perform basic tasks such as loading, sequencing, scheduling; and monitoring and control, a decision mechanism involving both operations staff and information systems that makes or confirms planning and control decisions. It is important that all these elements are effective in their own right and work together.

Are resource planning and control information integrated?

Resource planning and control involves vast amounts of information. Unless all relevant information is integrated it is difficult to make informed planning and control decisions. The most common method of doing this is through the use of integrated 'enterprise resource planning' (ERP) systems. These are information systems that have grown out of the more specialised and detailed material requirements planning (MRP) systems that have been common in the manufacturing sector for many years. MRP is treated in the supplement to this chapter. Investment in ERP systems often involves large amounts of capital and staff time. It also may mean a significant overhaul of the way the business organises itself. Not all investments in ERP have proved successful.

Are core planning and control activities effective?

Unless the resource planning and control system makes appropriate decisions at a detailed level, it cannot be effective. These detailed decisions fall into four overlapping categories. Loading is the activity of allocating work to individual processes or stages in the operation. Sequencing is the activity of deciding the order or priority in which a number of jobs will be processed. Scheduling is the activity of producing a detailed timetable showing when activities should start and end. Monitoring and control is the activity of detecting any deviation from what has been planned and acting to cope and replan as necessary. The theory of constraints (TOC) is a useful concept in resource planning and control, which emphasises the role of bottleneck stages or processes in planning and control.

DIAGNOSTIC QUESTION

Do resource planning and control have all the right elements?

Resource planning and control is concerned with managing the ongoing allocation of resources and activities to ensure that the operation's processes are both efficient and reflect customer demand for products and services. Planning and control activities are distinct but often overlap. Formally, planning determines what is *intended* to happen at some time in the future, while control is the process of *coping* when things do not happen as intended. Control makes the adjustments which help the operation to achieve the objectives that the plan has set, even when the assumptions on which the plan was based do not hold true.

Look at the resource planning and control activities in the following two organisations. One, Air France, is a very large and very complex network of operations and processes. The other, the service section of a BMW automotive dealership, is far smaller. However, although the challenges are different, the task of planning and controlling each operation's resources is surprisingly similar.

Operations control at Air France[1]

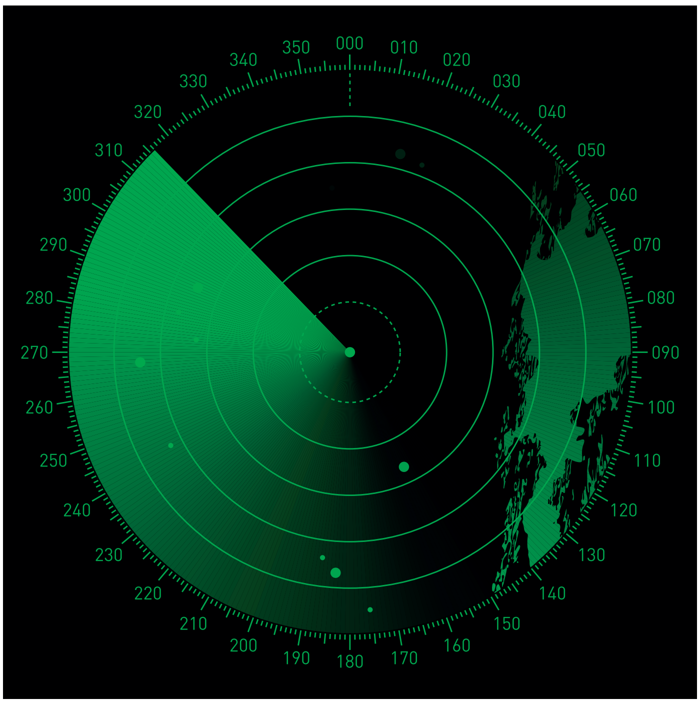

'In many ways a major airline can be viewed as one large planning problem which is usually approached as many independent, smaller (but still difficult) planning problems. The list of things which need planning seems endless: crews, reservation agents, luggage, flights, through trips, maintenance, gates, inventory, equipment purchases. Each planning problem has its own considerations, its own complexities, its own set of time horizons, its own objectives, but all are interrelated.' (Rikard Monet, Air France)

Air France has 80 flight planners working 24-hour shifts in their flight planning office at Roissy, Charles de Gaulle. Their job is to establish the optimum flight routes, anticipate any problems such as weather changes and minimise fuel consumption. Overall, the goals of the flight planning activity are first, and most important, safety followed by economy and passenger comfort. Increasingly powerful computer programs process the mountain of data necessary to plan the flights, but in the end many decisions still rely on human judgement. Even the most sophisticated expert systems only serve as support for the flight planners. Planning Air France's schedule is a massive job that includes the following.

- *Frequency* – for each airport how many separate services should the airline provide?
- *Fleet assignment* – which type of plane should be used on each leg of a flight?
- *Banks* – at any airline hub where passengers arrive and may transfer to other flights to continue their journey, airlines like to organise flights into 'banks' of several plans which arrive close together, pause to let passengers change planes and all depart close together.

- *Block times* – a block time is the elapsed time between a plane leaving the departure gate at an airport and arriving at its gate in the arrival airport. The longer the allowed block time, the more likely a plane will keep to schedule even if it suffers minor delays, but the fewer flights can be scheduled.
- *Planned maintenance* – any schedule must allow time for planes to have time at a maintenance base.
- *Crew planning* – pilot and cabin crew must be scheduled to allocate pilots to fly planes on which they are licensed and to keep within the maximum 'on duty' allowances.
- *Gate plotting* – if many planes are on the ground at the same time there may be problems in loading and unloading them simultaneously.
- *Recovery* – many things can cause deviations from any plan in the airline industry. Allowances must be built in that allow for recovery.

For flights within and between Air France's 12 geographic zones, the planners construct a flight plan that will form the basis of the actual flight only a few hours later. All planning documents need to be ready for the flight crew who arrive two hours before the scheduled departure time. Being responsible for passenger safety and comfort, the captain always has the final say and, when satisfied, co-signs the flight plan together with the planning officer.

Joanne manages the schedule[2]

Joanne Cheung is the Senior Service Adviser at a premier BMW dealership. Together with her team, she acts as the interface between customers who want their cars serviced and repaired, and the 16 technicians who carry out the work in their state-of-the-art workshop. *'There are 3 types of work that we have to organise,'* says Joanne, *'The first is performing repairs on customers' vehicles. They usually want this doing as soon as possible. The second type of job is routine servicing. It is usually not urgent so customers are generally willing to negotiate a time for this. The remainder of our work involves working on the pre-owned cars which our buyer has bought-in to sell on to customers. Before any of these cars can be sold, they have to undergo extensive checks. To some extent, we treat these categories of work slightly differently. We have to give good service to our internal car buyers, but there is some flexibility in planning these jobs. At the other extreme, emergency repair work for customers has to be fitted into our schedule as quickly as possible. If someone is desperate to have their car repaired at very short notice, we sometimes ask them to drop their car in as early as they can and pick it up as late as possible. This gives us the maximum amount of time to fit it into the schedule.'*

'There are a number of service options open to customers. We can book short jobs in for a fixed time and do it while they wait. Most commonly, we ask the customer to leave the car with us and collect it later. To help customers we have ten loan cars, which are booked out on a first-come first-served basis. Alternatively, the vehicle can be collected from the customer's home and delivered back there when it is ready. Our 4 drivers who do this are able to cope with up to 12 jobs a day.'

'Most days we deal with 50 to 80 jobs, taking from half-an-hour up to a whole day. To enter a job into our process all Service Advisers have access to the computer-based scheduling system. On-screen it shows the total capacity we have day-by-day, all the jobs that are booked in, the amount of free capacity still available, the number of loan cars available, and so on.

We use this to see when we have the capacity to book a customer in, and then enter all the customer's details. BMW have issued "standard times" for all the major jobs. However, you have to modify these standard times a bit to take account of circumstances. That is where the Service Adviser's experience comes in.'

'We keep all the most commonly used parts in stock, but if a repair needs a part that is not in stock, we can usually get it from the BMW parts distributors within a day. Every evening our planning system prints out the jobs to be done the next day and the parts that are likely to be needed for each job. This allows the parts staff to pick out the parts for each job so that the technicians can collect them first thing the next morning without any delay.'

'Every day we have to cope with the unexpected. A technician may find that extra work is needed, customers may want extra work doing, and technicians are sometimes ill, which reduces our capacity. Occasionally, parts may not be available so we have to arrange with the customer for the vehicle to be rebooked for a later time. Every day up to 4 or 5 customers just don't turn up. Usually, they have just forgotten to bring their car in so we have to rebook them in at a later time. We can cope with most of these uncertainties because our technicians are flexible in terms of the skills they have and are willing to work overtime when needed. Also, it is important to manage customer's expectations. If there is a chance that the vehicle may not be ready for them, it shouldn't come as a surprise when the try and collect it.'

What do they have in common?

The system set up by both Air France and the BMW dealership have a number of common elements. First, there is a type of acknowledgement that there should be an effective *customer interface* that translates the needs of customers into their implications for the operation. This involved setting the timetable of flights (frequency, timing, etc.) and the interfaces between flights (banks) in Air France. On a more individual scale Joanne needed to judge the degree of urgency of each job and feed back to the customer, managing their expectations where appropriate. Both planning and control systems also have a *supply interface* that translates the operation's plans in terms of the supply of parts, or fuel, ground services, crew availability, and so on. At the heart of each company's activities is a set of *core mechanics* that load capacity, prioritise, schedule, monitor and control the operation. The job of this decision-making is to reconcile the needs of the customers and the operation's resources in some way. For Joanne this involves attempting to maximise the utilisation of her workshop resources while keeping customers satisfied. Air France also has similar objectives, with customer comfort and safety being paramount. Also each operation is attempting some *information integration* that involves both computer-assisted information handling and the skills and experience of planning and control staff.

What elements should all resource planning and control have?

Figure 10.2 illustrates the elements that should be present in all planning and control systems. In more sophisticated systems they may even be extended to include the integration of this core operations resource planning and control task with other functional areas of the firm, such as finance, marketing and personnel. We deal with this cross-functional perspective when we discuss ERP later.

How does the system interface with customers?

The part of the resource planning and control system that manages the way customers' interact with the business on a day-to-day basis is called the 'customer interface' or sometimes 'demand management'. This is a set of activities that interface with both individual customers and the market more broadly. Depending on the business, these activities may include customer

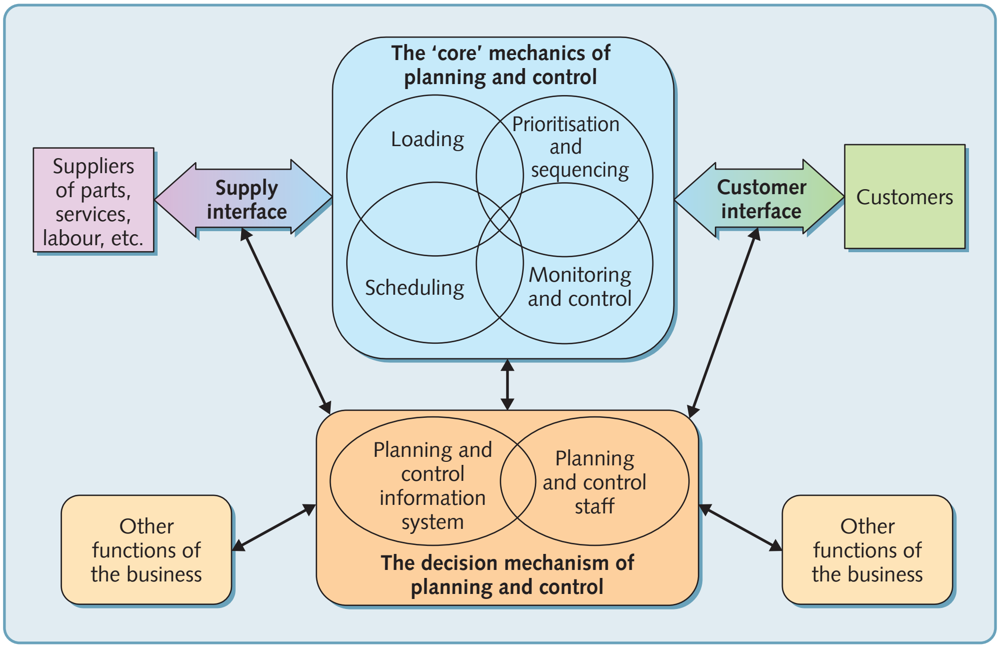

Figure 10.2 The key elements of a resource planning and control system

negotiation, order entry, demand forecasting, order promising, updating customers, keeping customer histories, post-delivery customer service and physical distribution.

Customer interface defines the customer experience

The customer interface is important because it defines the nature of the customer experience. It is the public face of the operation (the 'line of visibility' as it was called in Chapter 8). Therefore, it needs to be managed like any other 'customer processing' process, where the quality of the service, as the customer sees it, is defined by the gap between customers' expectations and their perceptions of the service they receive. Figure 10.3 illustrates a typical

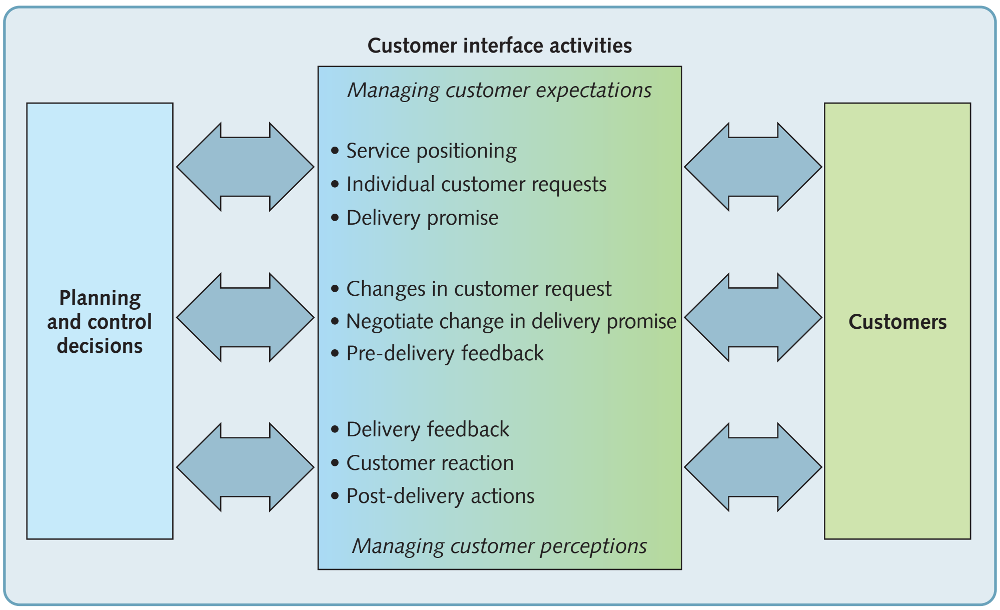

Figure 10.3 The customer interface as a 'customer experience'

customer experience of interacting with a planning and control customer interface. The experience itself will start before any customer contact is initiated. Customer expectations will have been influenced by the way the business presents itself through promotional activities, the ease with which channels of communication can be used (for example, design of the website), and so on. The question is, 'Does the communication channel give any indication of the kind of service response (for example, how long will we have to wait?) that the customer can expect?' At the first point of contact, when an individual customer requests services or products, their request must be understood, delivery possibly negotiated, and a delivery promise made. Prior to the delivery of the service or product, the customer may or may not change their mind, which in turn may or may not involve renegotiation delivery promises. Similarly, customers may require or value feedback on the progress of their request. At the point of delivery, not only are the products and services handed over to the customer, but there may also be an opportunity to explain the nature of the delivery and gauge customers' reactions. Following the completion of the delivery, there may also be a post-delivery action, such as a phone call to confirm that all is well.

OPERATIONS PRINCIPLE

Customers' perceptions of an operation will partially be shaped by the customer interface of its planning and control system.

As is usual with such customer experiences, managing customer expectations is particularly important in the early stages of the experience. For example, if there is a possibility that a delivery may be late (perhaps because of the nature of the service being requested) then that possibility is established as an element in the customer's expectations. As the experience continues, various interactions with the customer interface service to build up customer perceptions of the level of support and care exhibited by the operation.

The customer interface should reflect the operation's objectives

In managing a customer's experience, the customer interface element of the planning and control system is, in effect, operationalising the business's operations objectives. It may have to prioritise one type of customer over another. It may have to encourage some types of customer to transact business more than other (possibly less profitable) types of customer. It will almost certainly have to trade off elements of customer service against the efficiency and utilisation of the operations resources. No matter how sophisticated the customer interface technology, or how skilled the customer interface staff, this part of the planning and control system cannot operate effectively without clear priorities derived from the operation's strategic objectives.

The customer interface acts as a trigger function

Acceptance of an order should prompt the customer interface to trigger the operation's processes. Exactly what is triggered will depend on the nature of the business. For example, some building and construction companies, because they are willing to build almost any kind of construction, will keep relatively few of their own resources within the business, but rather hire them in when the nature of the job becomes evident. This is a 'resource-to-order' operation where the customer interface triggers the task of hiring in the relevant equipment (and possibly labour) and purchasing the appropriate materials. If the construction company confined itself to a narrower range of construction tasks, thereby making the nature of demand slightly more predictable, it would be likely to have its own equipment and labour permanently within the operation. Here, accepting a job would only need to trigger the purchase of the materials to be used in the construction, and the business is 'produce to order' operation. Some construction companies will construct pre-designed standard houses or apartments ahead of any firm demand for them. If demand is high, customers may place requests for houses before they are started or during their construction. In this case, the customer will form a backlog of demand and must wait. However the company is also taking the risk of holding a stock of unsold houses. Operations of this type 'produce-ahead-of-order'.

How does the system interface with suppliers?

The supplier interface provides the link between the activities of the operation itself and those of its suppliers. The timing and level of activities within the operation or process will have implications for the supply of products and services to the operation. Suppliers need to be informed so that they can make products and services available when needed. In effect, this is the mirror image of the customer interface. As such, the supplier interface is concerned with managing the supplier experience to ensure appropriate supply. Because the customer is not directly involved in this does not make it any less important. Ultimately, customer satisfaction will be influenced by supply effectiveness because that in turn influences delivery to customers.

Using the expectations–perception gap to judge the quality of the supplier interface function may at first seem strange. After all, suppliers are not customers as such. Yet, it is important to be a 'quality customer' to suppliers because this increases the chances of receiving high-quality service from them. This means that suppliers fully understand one's expectations because they have been made clear and unambiguous.

OPERATIONS PRINCIPLE

An operation's planning and control system can enhance or inhibit the ability of its suppliers to support delivery effectiveness.

The supplier interface has both a long- and short-term function. It must be able to cope with different types of long-term supplier relationship, and also handle individual transactions with suppliers. To do the former, it must understand the requirements of all the processes within the operation and also the capabilities of the suppliers (in large operations, there could be thousands of suppliers). Figure 10.4 shows a simplified sequence of events in the management of a typical supplier–operation interaction, which the supplier interface must facilitate. When the planning and control activity requests supply, the supplier interface must have identified potential suppliers and might also be able to suggest alternative materials or services if necessary. Formal requests for quotations may be sent to potential suppliers if no supply agreement exists. These requests might be sent to several suppliers or a smaller group, who may be 'preferred' suppliers. Just as it is important to manage customer expectations, it is important to manage supplier expectations, often prior to any formal supply of products or services. This issue was discussed in Chapter 7 as supplier development. To handle individual transactions, the supplier interface will need to issue formal purchase orders. These may be stand-alone documents or, more likely, electronic orders. Whatever the mechanisms, it is an important activity because it often forms

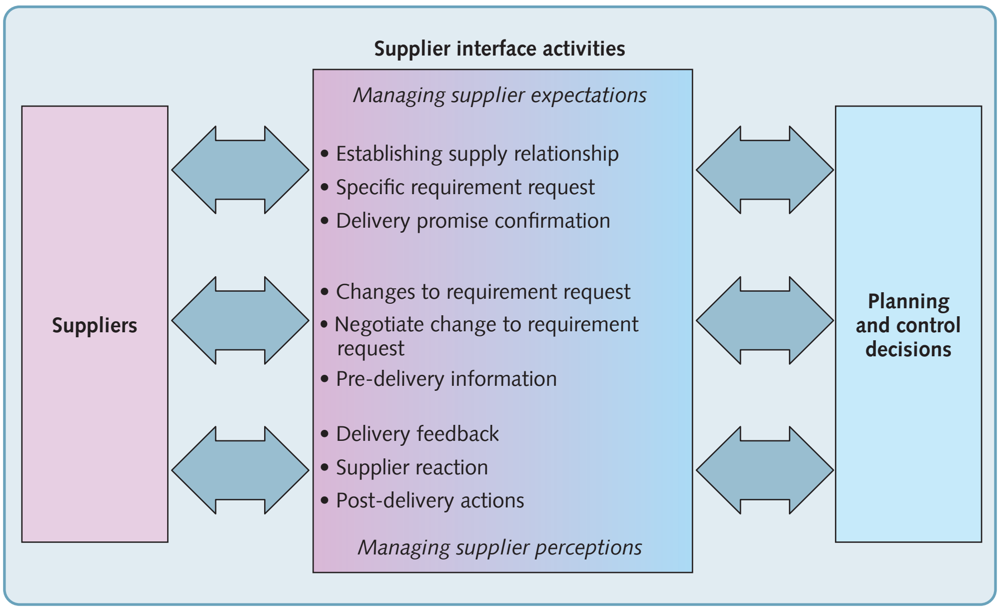

Figure 10.4 The supplier interface as a 'customer' experience

the legal basis of the contractual relationship between the operation and its supplier. Delivery promises will need to be formally confirmed. While waiting for delivery, it may be necessary to negotiate changes in supply, tracking progress to get early warning of potential changes to delivery. Also delivery supplier performance needs to be established and communicated with follow up as necessary.

How does the system perform basic planning and control calculations?

Resource planning and control requires the reconciliation of supply and demand in terms of the level and timing of activities within an operation or process. To do this, four overlapping activities are performed. These are loading, sequencing, scheduling, and monitoring and control. However, some caution is needed when using these terms. Different organisations may use them in different ways, and even textbooks in the area may adopt different definitions. Although these four activities are very closely interrelated, they do address different aspects of the resource planning and control task. Loading allocates tasks to resources in order to assess *what* level of activity will be expected of each part of the operation. Scheduling is more concerned with *when* the operation or process will do things. Sequencing is a more detailed set of decisions that determines *in what order* jobs pass through processes. Monitoring and control involves checking if *activities are going to* plan by observing what is actually happing in practice, and making adjustments as necessary. See Figure 10.5. This part of the planning and control system can be regarded as the engine room of the whole system insomuch as it calculates the consequences of planning and control decisions. Without understanding how these basic mechanisms work, it is difficult to understand how any operations are being planned and controlled. Because of their importance, we discuss the four interrelated activities later in the chapter.

Does the system integrate human with 'automated' decision-making?

Although computer-based resource planning and control systems are now widespread in many industries, much of the decision-making is still carried out partially by people. This is always likely to be the case because some elements of the task, such as negotiating with customers and suppliers, are difficult to automate. Yet the benefits of computer-aided decision-making are difficult to ignore. Unlike humans, computer-based planning and control can cope with immense complexity, both in terms of being able to model the interrelationship between decisions and in terms of being able to store large quantities of information. However, humans are generally

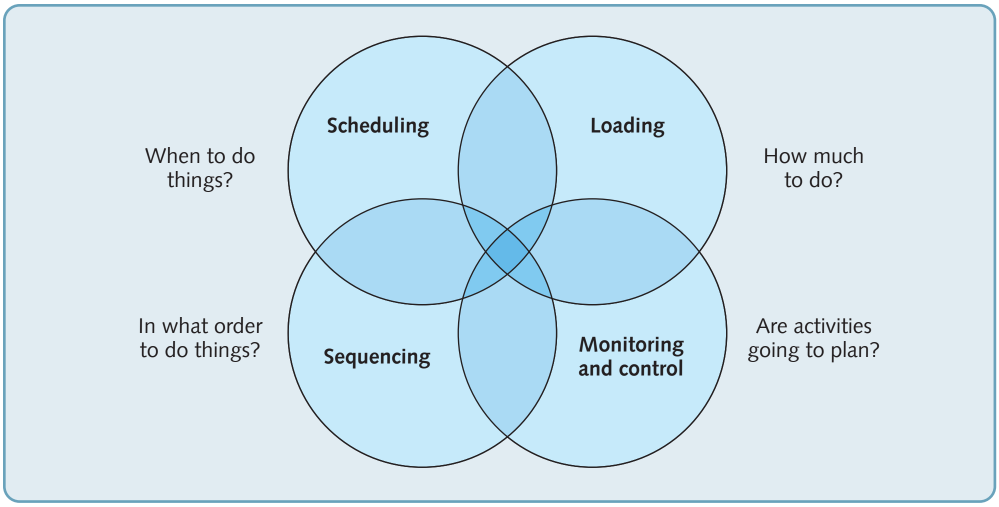

Figure 10.5 **The 'core mechanisms' of planning and control**

better at many of the 'soft' qualitative tasks that can be important in planning and control. In particular, humans are good at the following:

- *Flexibility, adaptability and learning.* Humans can cope with ambiguous, incomplete, inconsistent and redundant goals and constraints. In particular, they can deal with the fact that planning and control objectives and constraints may not be stable for longer than a few hours.
- *Communication and negotiation.* Humans are able to understand and sometimes influence the variability inherent in an operation. They can influence job priorities and sometimes processing times. They can negotiate between internal processes and communicate with customers and suppliers in a way that could minimise misunderstanding.
- *Intuition.* Humans can fill in the blanks of missing information that is required to plan and control. They can accumulate the tacit knowledge about what is, and what may be, really happening with the operation's processes.

These strengths of human decision-making versus computer decision-making provide a clue as to what should be the appropriate degree of automation built into decision-making in this area. When planning and controlling stable and relatively straightforward processes that are well understood, decision-making can be automated to a greater degree than processes that are complex, unstable and poorly understood.

DIAGNOSTIC QUESTION

Are resource planning and control information integrated?

One of the most important issues in resource planning and control is managing the sometimes vast amounts of information generated, not just from the operations function, but from almost every other function of the business. Unless all relevant information is brought together and integrated it is difficult to make informed planning and control decisions. This is what Enterprise Resource Planning (ERP) is about. It has been defined as, *a complete enterprise wide business solution. The ERP system consists of software support modules such as: marketing and sales, field service, product design and development, production and inventory control, procurement, distribution, industrial facilities management, process design and development, manufacturing, quality, human resources, finance and accounting, and information services. Integration between the modules is stressed without the duplication of information.*

OPERATIONS PRINCIPLE

Planning and control systems should integrate information from all relevant organisational functions.

The origins of ERP

Enterprise resource planning has spawned a huge industry devoted to developing the computer systems needed to drive it. The (now) large companies which have grown almost exclusively on the basis of providing ERP systems include SAP, Oracle and Baan. Yet ERP is the one of the latest (and most important) stages in a development that started with materials requirements planning(MRP), an approach that became popular during the 1970s, although the planning and control logic that underlies it had been known for some time. It is a method (simple in principle but complex in execution) of translating a statement of required output into a plan for all the activities that must take place to achieve the required output. What popularised MRP was the availability of computer power to drive the basic planning and control mathematics in a fast, efficient and, most importantly, flexible manner. MRP is treated in the supplement to

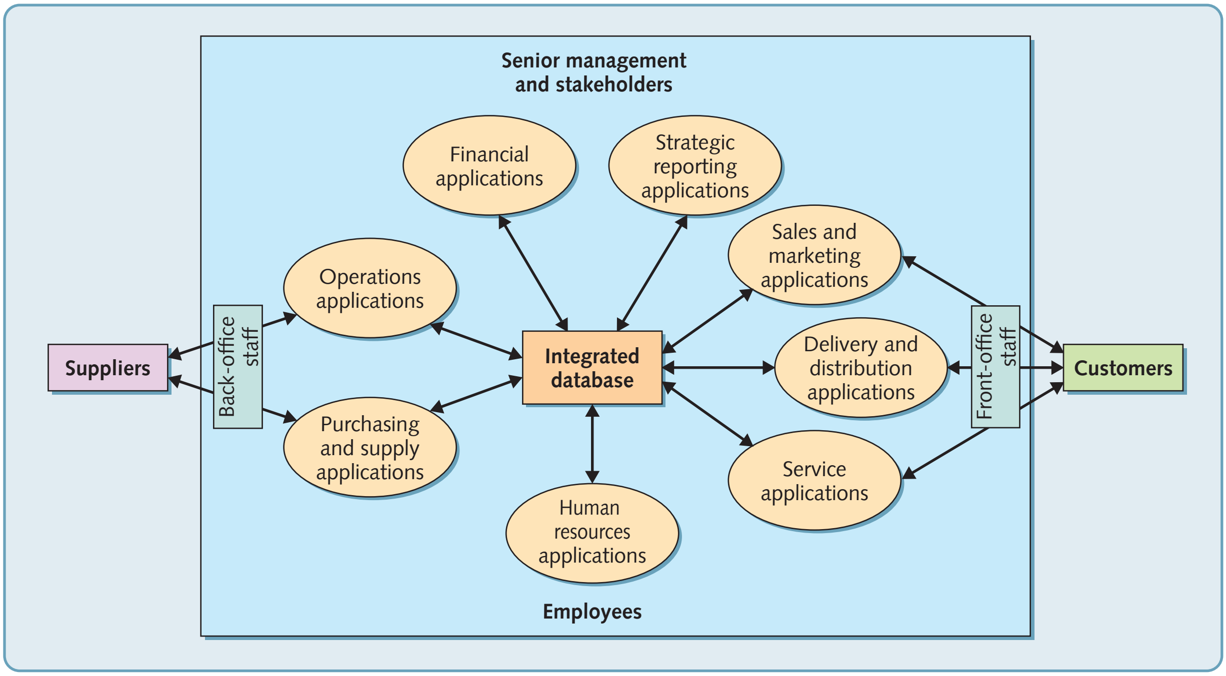

Figure 10.6 ERP integrates information from all parts of the organisation

this chapter. Manufacturing resource planning (MRP II) expanded out of MRP during the 1980s. This extended concept has been described as a game plan for planning and monitoring all the resources of a manufacturing company: manufacturing, marketing, finance and engineering. Again, it was a technology innovation that allowed the development. Local area networks (LANs), together with increasingly powerful desktop computers, allowed a much higher degree of processing power and communication between different parts of a business.

The strength of MRP and MRP II lay always in the fact that it could explore the *consequences* of any changes to what an operation was required to do. So, if demand changed, the MRP system would calculate all the 'knock-on' effects and issue instructions accordingly. The same principle applies to ERP, but on a much wider basis. ERP systems allow decisions and databases from all parts of the organisation to be integrated so that the consequences of decisions in one part of the organisation are reflected in the planning and control systems of the rest of the organisation (see Figure 10.6).

ERP changes the way companies do business

Arguably the most significant issue in many company's decision to buy an off-the-shelf ERP system is that of its compatibility with the company's current business processes and practices. Experience of ERP installation suggests that it is extremely important to make sure that one's current way of doing business will fit (or can be changed to fit) with a standard ERP package. One of the most common reasons for not installing ERP is incompatibility between the assumptions in the software and the operating practice of core business processes. If a business's current processes do not fit, they can either change their processes to fit the ERP package, or modify the software within the ERP package to fit their processes.

> **OPERATIONS PRINCIPLE**
>
> ***ERP systems are only fully effective if the way a business organises its processes is aligned with the underlying assumption of its ERP system.***

Both of these options involve costs and risks. Changing business practices that are working well will involve reorganisation costs as well introducing the potential for errors to creep into the processes. Adapting the software will both slow down the project and introduce potentially dangerous software 'bugs' into the system. It would also make it difficult to upgrade the software later on.

ERP installation can be particularly expensive. Attempting to get new systems and databases to talk to old (sometimes called *legacy*) systems can be very problematic. Not surprisingly, many companies choose to replace most, if not all, their existing systems simultaneously. New common systems and relational databases help to ensure the smooth transfer of data between different parts of the organisation.

In addition to the integration of systems, ERP usually includes other features which make it a powerful planning and control tool:

- It is based on a client/server architecture; i.e., access to the information systems is open to anyone whose computer is linked to central computers.
- It can include decision support facilities that enable operations decision makers to include the latest company information.
- It is often linked to external extranet systems, such as the electronic data interchange (EDI) systems, which are linked to the company's supply chain partners.
- It can be interfaced with standard applications programs which are in common use by most managers, such as spreadsheets, etc.
- Often, ERP systems are able to operate on most common platforms such as Windows NT or UNIX, or Linux.

The benefits of ERP

ERP is generally seen as having the potential to significantly improve the performance of many companies in many different sectors. This is partly because of the very much enhanced visibility that information integration gives, but it is also a function of the discipline that ERP demands. Yet this discipline is itself a 'double-edged' sword. On one hand, it 'sharpens up' the management of every process within an organisation, allowing best practice (or at least common practice) to be implemented uniformly through the business. No longer will individual idiosyncratic behaviour by one part of a company's operations cause disruption to all other processes. On the other hand, it is the rigidity of this discipline that is both difficult to achieve and (arguably) inappropriate for all parts of the business. Nevertheless, the generally accepted benefits of ERP are as follows:

- Greater visibility of what is happening in all parts of the business.
- Forcing the business process-based changes that potentially make all parts of the business more efficient.
- Improved control of operations that encourages continuous improvement (albeit within the confines of the common process structures).
- More sophisticated communication with customers, suppliers, and other business partners, often giving more accurate and timely information.
- Integrating whole supply chains including suppliers' suppliers and customers' customers.

Web-integrated ERP

An important justification for embarking on ERP is the potential it gives to link up with the outside world; for example, by integrating its external communication systems into its internal ERP systems. However, as has been pointed out by some critics of the ERP software companies, different types of external company often need different types of information. Customers need to check the progress of their orders and invoicing, whereas suppliers and other partners want access to the details of operations planning and control. Not only that, but they want access all the time. The internet is always there, but ERP systems are often complex and need periodic maintenance. This can mean that every time the ERP system is taken offline for routine maintenance or other changes, the web site also goes offline. To combat this, some companies configure their ERP and e-commerce links in such a way that they can be decoupled so that ERP can be periodically shut down without affecting the company's web presence.

Supply network ERP

The step beyond integrating internal ERP systems with immediate customers and suppliers is to integrate it with the systems of other businesses throughout the supply network. This is often exceptionally complicated. Not only do different ERP systems have to communicate together, they have to integrate with other types of system. For example, sales and marketing functions often use systems such as customer relationship management (CRM) systems that manage the complexities of customer requirements, promises and transactions. Getting ERP and CRM systems to work together is itself often difficult. Nevertheless, such web-integrated ERP, or 'c-commerce' (collaborative commerce) applications are emerging and starting to make an impact on the way companies do business. Although a formidable task, the benefits are potentially great. The costs of communicating between supply network partners could be dramatically reduced and the potential for avoiding errors as information and products move between partners in the supply chain are significant. Yet such transparency also brings risks. If the ERP system of one operation within a supply chain fails for some reason, it may block the effective operation of the whole integrated information system throughout the network.

EXAMPLE

What a waste[3]

Not only can ERP implementation go wrong, even when undertaken by experienced professionals, sometimes it can end up in the law courts. Waste ManagementInc. is the leading provider of waste and environmental services in North America. So when it announced that it was suing SAP (see earlier 'operations in practice' example) over the failure of an ERP implementation, planning and control systems practitioners took notice. Waste Management said that it was seeking the recovery of more than $100 million in project expenses as well as 'the savings and benefits that the SAP software was promised to deliver to Waste Management.' It said that SAP promised that the software could be fully implemented throughout all of Waste Management within 18 months, and that its software was an 'out-of-the-box' solution that would meet Waste Management's needs without any customisation or enhancements,'

Waste Management signed a sales pact with SAP, but according to Waste Management, '*Almost immediately following execution of the agreements, the SAP implementation team discovered significant "gaps" between the software's functionality and Waste Management's business requirements. Waste Management has discovered that these gaps were already known to the product development team in Germany even before the SLA (service level agreement) was signed*.' But members of SAP's implementation team had reportedly blamed Waste Management for the functional gaps and had submitted change orders requiring that Waste Management pay for fixing them. Five years later, the dispute was settled when SAP made a one-time cash payment to Waste Management.

Sales and operations planning (S&OP)

One of the problems with traditional operations planning and control is that, although several functions were often routinely involved in the process, each function could have a very different set of objectives. For example, marketing could be interested in maximising revenues and ensuring continuity of delivery to customers. Operations are likely to be under pressure to minimise cost (perhaps achieved through relatively long and stable operating levels). Finance will be interested in reducing working capital and inventory and reducing fixed costs. And so on. Yet these, and other functions such as engineering or human resources management, are all impacted by operations planning decisions and are probably involved in their own planning processes that partly depend on the output from the operations planning process. S&OP was first promoted as an important element of planning as manufacturing resource planning (see the supplement to this chapter) became a commonly used process. Early manufacturing resource planning implementations were often made less effective by the system being driven by unachievable plans. This is the dilemma that S&OP is intended to address. It is a planning process that attempts to ensure that all tactical plans are aligned across the business's various functions and with the company's longer-term strategic plans.

It is a formal business process that looks over a period of 18 to 24 months ahead. In other words, it is not a purely short-term process. In fact S&OP developed as an attempt to integrate short- and longer-term planning, as well as integrating the planning activities of key functions. It is an aggregated process that does not deal with detailed activities, but rather focuses on the overall (often aggregated) volume of output. Generally, it is a process that happens monthly, and tends to take place at a higher level, involving more senior management than traditional operations planning. S&OP also goes by many names. It can be called integrated business planning, integrated business management, integrated performance management, rolling business planning, and regional business management, to name a few. It has also been noted[4] that some organisations continue to use the phrase 'S&OP', although they may mean something quite different.

DIAGNOSTIC QUESTION

Are core planning and control activities effective?

All resource planning and control activity eventually relies on a set of calculations that guide how much work to load onto different parts of the operation, when different activities should be performed, in what order individual jobs should be done and how processes can be adjusted if they have deviated from plan. These calculations can be thought of as the 'engine room' of the whole resource planning and control system. Although the algorithms that guide the calculations are often embedded within computer-based systems, it is worthwhile understanding some of the core ideas on which they are based. These fall into four overlapping categories: loading, scheduling, sequencing, and monitoring and control.

Loading

Loading is the amount of work that is allocated to a process or stage or whole process. It is a capacity related issue that will attempt to reconcile how much the operation or the process is expected to do with how much the operation or process can do. Essentially, the loading activity calculates the consequences on individual parts of the operation of the operation's overall workload. It may or may not assume realistic capacity limits on what can be loaded. If it does, it is called finite loading;

OPERATIONS PRINCIPLE

For any given level of demand a planning and control system should be able to indicate the implications for the loading on any part of the operation.

if not, it is called infinite loading. Finite loading is an approach that only allocates work to a work centre (a person, a machine, or perhaps a group of people or machines) up to a set limit. This limit is the estimate of capacity for the work centre (based on the times available for loading). Work over and above this capacity is not accepted. Figure 10.7(a) shows that the load on the work centre is not allowed to exceed the capacity limit. Finite loading is particularly relevant for operations where:

- *it is possible to limit the load* – for example, it is possible to run an appointment system for a general medical practice or a hairdresser
- *it is necessary to limit the load* – for example, for safety reasons only a finite number of people and weight of luggage are allowed on aircraft
- *the cost of limiting the load is not prohibitive* – for example, the cost of maintaining a finite order book at a specialist sports car manufacturer does not adversely affect demand, and may even enhance it.

Infinite loading is an approach to loading work that does not limit accepting work, but instead tries to cope with it. Figure 10.7(b) illustrates a loading pattern where capacity constraints have not been used to limit loading. Infinite loading is relevant for operations where:

- *it is not possible to limit the load* – for example, an accident and emergency department in a hospital should not turn away arrivals needing attention
- *it is not necessary to limit the load* – for example, fast-food outlets are designed to flex capacity up and down to cope with varying arrival rates of customers. During busy periods, customers accept that they must queue for some time before being served. Unless this is extreme, the customers might not go elsewhere
- *the cost of limiting the load is prohibitive* – for example, if a retail bank turned away customers at the door because a set amount were inside, customers would feel less than happy with the service.

In complex planning and control activities where there are multiple stages, each with different capacities and with a varying mix arriving at the facilities, such as a machine shop in an engineering company, the constraints imposed by finite loading make loading calculations complex and not worth the considerable computational power that would be needed.

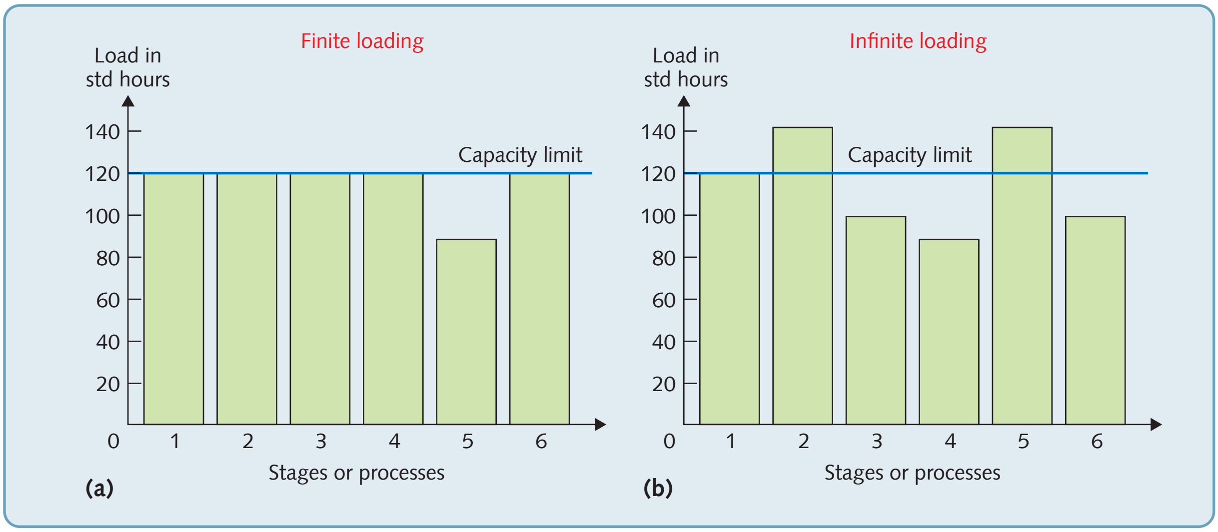

Figure 10.7 Finite and infinite loading

Sequencing

After the 'loading' of work onto processes, the order or sequence in which it will be worked on needs to be determined. This task is called 'sequencing'. The priorities given to work in an operation are often determined by some predefined set of sequencing rules. Some of these are summarised below.

Customer priority

This allows an important or aggrieved customer, or item, to be prioritised irrespective of their order of arrival. Some banks, for example, give priority to important customers. Accident and Emergency departments in hospitals must rapidly devise a schedule that prioritises patients presenting symptoms of a serious illness. Hospitals have developed 'triage systems', whereby medical staff hurriedly sort through the patients to determine their relative urgency.

Due date (DD)

Work is sequenced according to when it is 'due' for delivery, irrespective of the size of each job or the importance of each customer. For example, a support service in an office block, such as a reprographic unit, may sequence the work according to when the job is needed. Due date sequencing usually improves delivery reliability and average delivery speed, but may not provide optimal productivity.

Last in first out (LIFO)

This is usually selected for practical reasons. For example, unloading an elevator is more convenient on a LIFO basis, as there is only one entrance and exit. LIFO has a very adverse effect on delivery speed and reliability.

First in first out (FIFO)

Also called 'first-come, first-served' (FCFS), it is a simple and equitable rule, used especially when queues are evident to customers, as in theme parks.

Longest operation time first (LOT)

Executing the longest job first has the advantage of utilising work centres for long periods, but, although utilisation may be high (therefore cost relatively low), this rule does not take into account delivery speed, delivery reliability or flexibility.

Shortest operation time first (SOT)

This sends small jobs quickly through the process, so achieving output quickly, and enabling revenue to be generated quickly. Short-term delivery performance may be improved, but productivity and the throughput time of large jobs is likely to be poor.

WORKED EXAMPLE

Steve Smith is a website designer. Returning from his annual vacation (he finished all outstanding jobs before he left), five design jobs are given to him upon arrival at work. He gives them the codes A to E. Steve has to decide in which sequence to undertake the jobs. He wants both to minimise the average time the jobs are tied up in his office and, if possible, to meet the deadlines (delivery times) allocated to each job.

His first thought is to do the jobs in the order they were given to him, that is, first in first out (FIFO).

Sequencing rule: first in first out (FIFO)

Sequence of jobs	*Process time (days)*	*Start time*	*Finish time*	*Due date*	*Lateness (days)*
A	5	0	5	6	0
B	3	5	8	5	3
C	6	8	14	8	6
D	2	14	16	7	9
E	1	16	17	3	14
Total time in process		60	**Total lateness**		32
Average time in process (total/5)		12	**Average lateness (total/5)**		6.4

Alarmed by the average lateness, he tries the due date (DD) rule.

Sequencing rule: due date (DD)

Sequence of jobs	*Process time (days)*	*Start time*	*Finish time*	*Due date*	*Lateness (days)*
E	1	0	1	3	0
B	3	1	4	5	0
A	5	4	9	6	3
D	2	9	11	7	4
C	6	11	17	8	9
Total time in process		42	**Total lateness**		16
Average time in process (total/5)		8.4	**Average lateness (total/5)**		3.2

Better! But Steve tries out the shortest operation time (SOT) rule.

Sequencing rule: shortest operation time (SOT)

Sequence of jobs	*Process time (days)*	*Start time*	*Finish time*	*Due date*	*Lateness (days)*
E	1	0	1	3	0
D	2	1	3	7	0
B	3	3	6	5	1
A	5	6	11	6	5
C	6	11	17	8	9
Total time in process		38	**Total lateness**		16
Average time in process (total/5)		7.6	**Average lateness (total/5)**		3.2

This gives the same degree of average lateness but with a lower average time in the process. Steve decides to use the SOT rule.

There are many sequencing rules of differing complexity. Although different rules will perform differently depending on the circumstances of the sequencing problem, in practice the SOT rule generally performs well.

EXAMPLE

Can airline passengers be sequenced?

Like many before him, Dr Jason Steffen, a professional astrophysicist from the world famous Fermilab was frustrated by the time it took to load him and his fellow passengers onto the aircraft. He decided to devise a way to make the experience a little less tedious. So, for a while, he neglected his usual work of examining extra-solar planets, dark matter and cosmology, and experimentally tested a faster method of boarding aircraft. He found that by changing the sequence in which passengers are loaded onto the aircraft, airlines could potentially save both time and money. Using a computer simulation and the arithmetic techniques routinely used in his day-to-day job, he found what seemed to be a superior sequencing method. In fact, the most common way of boarding passenger planes proved to be the least efficient. This is called the 'block method' where blocks of seats are called for boarding, starting from the back. Previously other experts in the airline industry had suggested boarding those in window seats first, followed by middle and aisle seats. This is called the Wilma method. But according to Dr Steffen's simulations, two things slow down the boarding process. The first is that passengers may be required to wait in the aisle while those ahead of them store their luggage before they can take their seat. The second is that passengers already seated in aisle or middle seats frequently have to rise and move into the aisle to let others take seats nearer the window. So Dr Steffen's suggested a variant of the Wilma method that minimised the first type of disturbance and eliminated the second. He suggested boarding in alternate rows, progressing from the rear forward, window seats first. Using this approach (now called the Steffen method) first, the window seats for every other row on one side of the plane are boarded. Next, alternate rows of window seats on the opposite side are boarded. Then, the window seats in the skipped rows are filled in on each side. The procedure then repeats with the middle seats and the aisles. See Figure 10.8.

Later, the effectiveness of the various approached was tested using a mock-up of a Boeing 757 aircraft and 72 luggage-carrying volunteers. Five different scenarios were tested: 'block' boarding in groups of rows from back to front; one by one from back to front; the 'Wilma method'; the Steffen method; and completely random boarding. In all cases, parent–child pairs were allowed to board first. It was assumed that families were likely to want to stay together. As Dr Steffen had predicted, the conventional block approach came out as the slowest, with the strict back-to-front approach not much better. Completely random boarding (unallocated seating), which is used by several low-cost airlines fared much better, most probably because it randomly avoids space conflicts. The times for fully boarding the 70 passengers using each method were as follows: 'Block' boarding – 6:54 min; Back-to-front – 6:11 min; Random boarding – 4:44 min; Wilma method – 4:13 min; Steffen method – 3:36 min.

The big question is, 'would passengers really be prepared to be sequenced in this way as they queue to board the aircraft?' Some airlines argue that directing passengers on to a plane is a little like herding cats. But if they could adopt Dr Steffen's system it would save time for customers and very significant amounts of money for airlines.

Scheduling

Scheduling is the activity of producing a detailed timetable showing when activities should start and end. Schedules are familiar in many consumer environments; for example, a bus schedule that shows the time each bus is due to arrive at each stage of the route. But, although

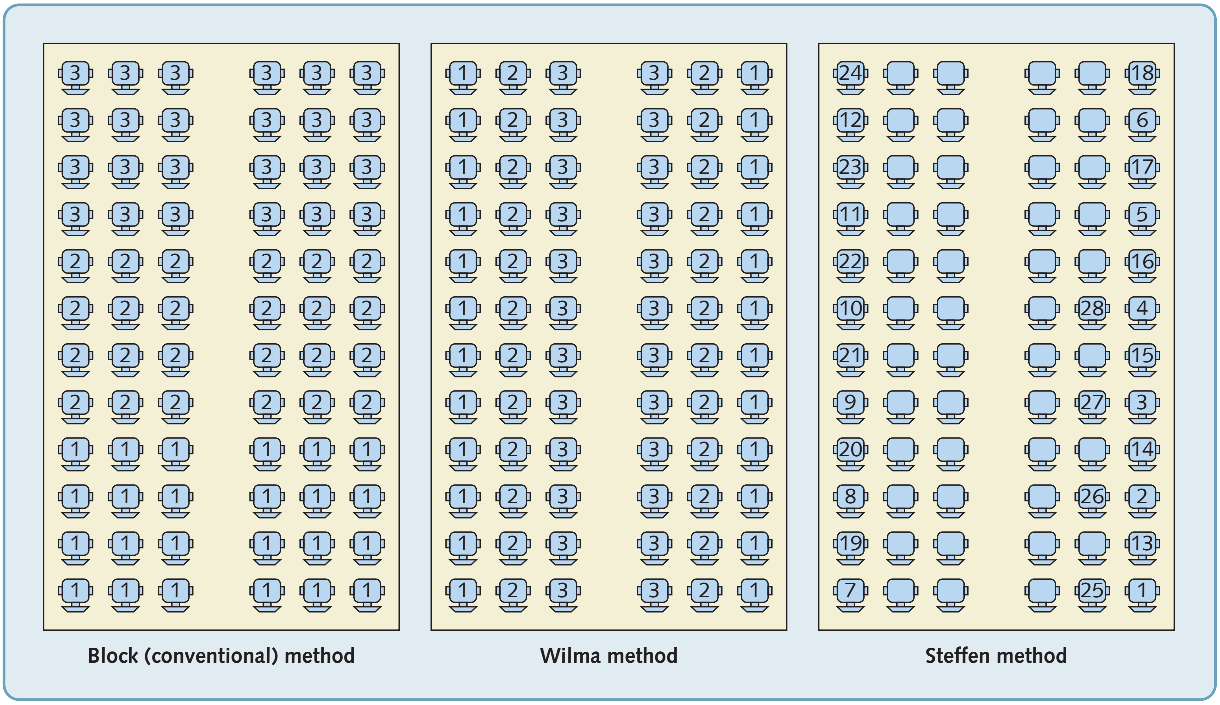

Figure 10.8 **The best way to sequence passengers onto an aircraft**

familiar, it is one of the most complex tasks in operations and process management. Schedules may have to deal simultaneously with many activities and several different types of resource, probably with different capabilities and capacities. Also the number of possible schedules increases rapidly as the number of activities and resources increase. If one process has five different jobs to process, any of the five jobs could be processed first and, following that, any one of the remaining four jobs, and so on. This means that there are: $5 \times 4 \times 3 \times 2 = 120$ different schedules possible. More generally, for n jobs there are $n!$ (factorial n, or $n \times (n - 1) \times (n - 2) \ldots \times 1$))different ways of scheduling the jobs through a single process or stage. If there is more than one process or stage, there are $(n!)^m$ possible schedules. Where n is the number of jobs and m is the number of process or stage. In practical terms, this means that there are often many millions of feasible schedules, even for relatively small operations. This is why scheduling rarely attempts to provide an 'optimal' solution but rather satisfies itself with an 'acceptable' feasible one.

OPERATIONS PRINCIPLE

An operation's planning and control system should allow for the effects of alternative schedules to be assessed.

Gantt charts

The most common method of scheduling is by use of the Gantt chart. This is a simple device which represents time as a bar, or channel, on a chart. The start and finish times for activities can be indicated on the chart and sometimes the actual progress of the job is also indicated. The advantages of Gantt charts are that they provide a simple visual representation both of what should be happening and of what actually is happening in the operation. Furthermore, they can be used to 'test out' alternative schedules. It is a relatively simple task to represent alternative schedules (even if it is a far from simple task to find a schedule which fits all the resources satisfactorily). Figure 10.9 illustrates a Gantt chart for a specialist software developer. It indicates the progress of several jobs as they are expected to progress through five stages of the process. Gantt charts are not an optimising tool, they merely facilitate the development of alternative schedules by communicating them effectively.

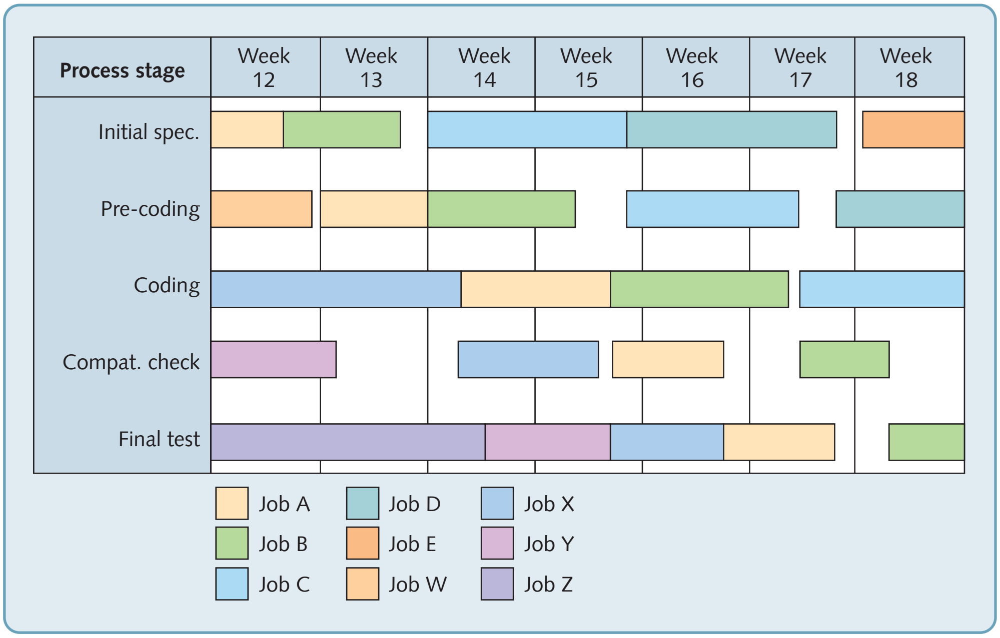

Figure 10.9 Gantt chart showing the schedule for jobs at each process stage

Scheduling work patterns

Where the dominant resource in an operation is its staff, then the schedule of work times effectively determines the capacity of the operation itself. Scheduling needs to make sure enough people are working at any time to provide a capacity appropriate for the level of demand. Operations such as call centres and hospitals, which must respond directly to customer demand, will need to schedule the working hours of their staff with demand in mind. For example, Figure 10.10 shows the scheduling of shifts for a small technical 'hot line' technical support service for the software company. Its service times are 04.00 to 20.00 on Monday, 04.00 to 22.00 from Tuesday to Friday, 06.00 to 22.00 on Saturday, and 10.00 to 20.00 on Sunday. Demand is heaviest from Tuesday to Thursday, starts to decrease on Friday, is low over the weekend and starts to increase again on Monday. The scheduling task for this type of problem can be considered over different time scales, two of which are shown in Figure 10.10. During

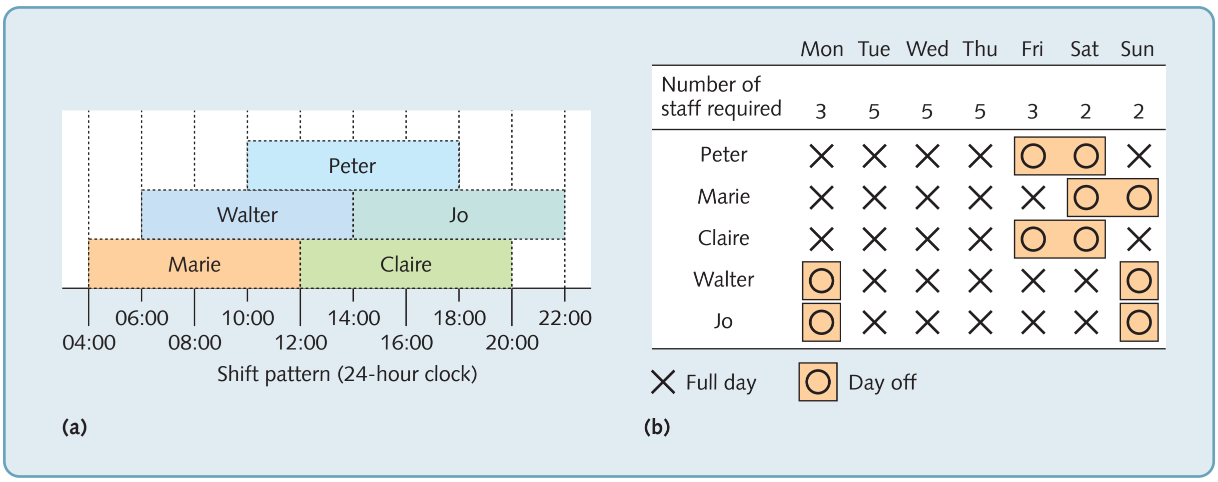

Figure 10.10 Shift allocation for the technical 'hot line' (a) on a daily basis, (b) on a weekly basis

the day, working hours need to be agreed with individual staff members. During the week, days off need to be agreed. During the year, vacations, training periods and other blocks of time where staff are unavailable need to be agreed. All this has to be scheduled such that:

- capacity matches demand
- the length of each shift is neither excessively long, nor too short to be attractive to staff
- working at unsocial hours is minimised
- days off match agreed staff conditions (for example); in this example, staff prefer two consecutive days off every week
- vacation and other 'time-off' blocks are accommodated
- sufficient flexibility is built into the schedule to cover for unexpected changes in supply (staff illness) and demand (surge in customer calls).

Scheduling staff times is one of the most complex of scheduling problems. In the relatively simple example shown in Figure 10.10 we have assumed that all staff have the same level and type of skill. In very large operations with many types of skill to schedule and uncertain demand (for example a large hospital) the scheduling problem becomes extremely complex. Some mathematical techniques are available but most scheduling of this type is, in practice, solved using heuristics (rules of thumb), some of which are incorporated into commercially available software package.

Theory of constraints (TOC)

An important concept, closely related to scheduling that recognises the importance of planning to known capacity constraints is the theory of constraints (TOC). It focuses scheduling effort on the bottleneck parts of the operation. By identifying the location of constraints, working to remove them, and then looking for the next constraint, an operation is always focusing on the part that critically determines the pace of output. The approach which uses this idea is called optimised production technology (OPT). Its development and the marketing of it as a proprietary software product were originated by Eliyahu Goldratt.[5] It helps to schedule production systems to the pace dictated by the most heavily loaded resources, that is, bottlenecks. If the rate of activity in any part of the system exceeds that of the bottleneck, then items are being produced that cannot be used. If the rate of working falls below the pace at the bottleneck, then the entire system is underutilised. The 'principles' of underlying OPT demonstrate this focus on bottlenecks.

OPT principles

1. Balance flow, not capacity. It is more important to reduce throughput time rather than achieving a notional capacity balance between stages or processes.
2. The level of utilisation of a non-bottleneck is determined by some other constraint in the system, not by its own capacity. This applies to stages in a process, processes in an operation, and operations in a supply network.
3. Utilisation and activation of a resource are not the same. According to the TOC a resource is being *utilised* only if it contributes to the entire process or operation creating more output. A process or stage can be *activated* in the sense that it is working, but it may only be creating stock or performing other non-value added activity.
4. An hour lost (not used) at a bottleneck is an hour lost forever out of the entire system. The bottleneck limits the output from the entire process or operation, therefore the underutilisation of a bottleneck affects the entire process or operation.
5. An hour saved at a non-bottleneck is a mirage. Non-bottlenecks have spare capacity anyway. Why bother making them even less utilised?
6. Bottlenecks govern both throughput and inventory in the system. If bottlenecks govern flow, then they govern throughput time, which in turn governs inventory.

7. You do not have to transfer batches in the same quantities as you produce them. Flow will probably be improved by dividing large production batches into smaller ones for moving through a process.
8. The size of the process batch should be variable, not fixed. Again, from the EBQ model, the circumstances that control batch size may vary between different products.
9. Fluctuations in connected and sequence-dependent processes add to each other rather than averaging out. So, if two parallel processes or stages are capable of a particular average output rate, in parallel, they will never be able to produce the same average output rate.
10. Schedules should be established by looking at all constraints simultaneously. Because of bottlenecks and constraints within complex systems, it is difficult to work out schedules according to a simple system of rules. Rather, all constraints need to be considered together.

Monitoring and control

Having created a plan for the operation through loading, sequencing and scheduling, each part of the operation has to be monitored to ensure that activities are happening as planned. Any deviation can be rectified through some kind of intervention in the operation, which itself will probably involve some replanning. Figure 10.11 illustrates a simple view of control. The output from a work centre is monitored and compared with the plan that indicates what the work centre is supposed to be doing. Deviations from this plan are taken into account through a replanning activity and the necessary interventions made (on a timely basis) to the work centre, which will ensure that the new plan is carried out. Eventually, however, some further deviation from planned activity will be detected and the cycle is repeated.

> OPERATIONS PRINCIPLE
>
> ***A planning and control system should be able to detect deviations from plans within a timescale that allows an appropriate response.***

Push and pull control

An element of control is periodic intervention into processes and operations. A key distinction is between intervention signals that push work through processes and operations and those that pull work only when it is required. In a *push* system of control, activities are scheduled by means of a central system and completed in line with central instructions, such as an MRP system (see the supplement to this chapter). Each work centre pushes out work without considering whether the succeeding work centre can make use of it. Deviations from plan are noted by the central operations planning and control system, and plans adjusted as required. In a *pull* system of control, the pace and specification of what is done are set by the succeeding 'customer' workstation, which 'pulls' work from the preceding (supplier) workstation. The customer acts as the only 'trigger' for movement. If a request is not passed

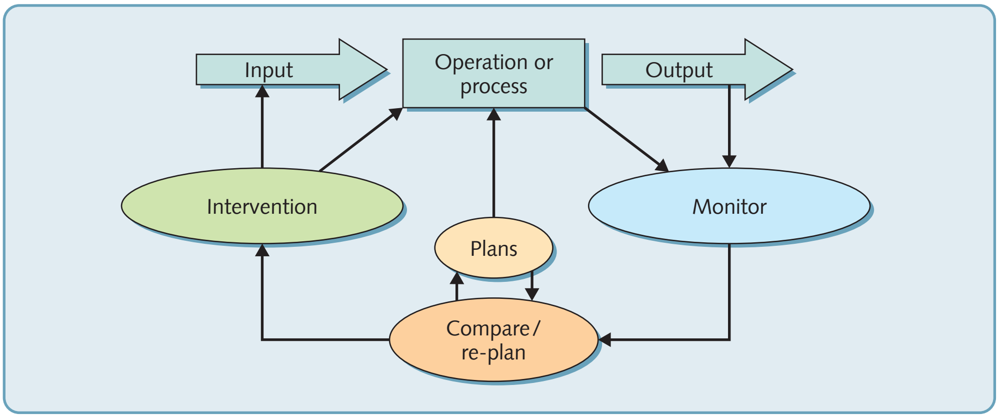

Figure 10.11 A simple model of control

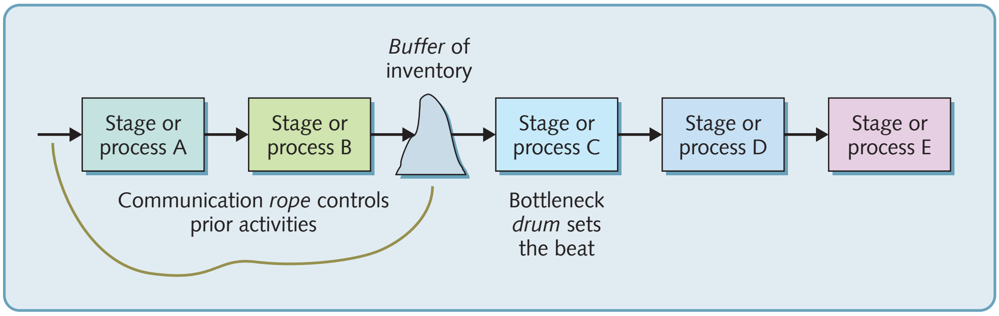

Figure 10.12 **The drum, buffer, rope concept**

back from the customer to the supplier, the supplier cannot produce anything or move any materials. A request from a customer not only triggers production at the supplying stage, but also prompts the supplying stage to request a further delivery from its own suppliers. In this way, demand is transmitted back through the stages from the original point of demand by the original customer.

Push systems of control are more formal and require significant decision-making or computing power when it is necessary to replan in the light of events. But push control can cope with very significant changes in circumstances such as major shifts in output level or product mix. By contrast, pull control is more self-adjusting in the sense that the nature rules that govern relationships between stages or processes cope with deviations from plan without reference to any higher decision-making authority. But there are limits to the extent that this can cope with major fluctuations in demand. Pull control works best when conditions are relatively stable. Understanding the differences between push and pull is also important because they have different effects in terms of their propensities to accumulate inventory. Pull systems are far less likely to result in inventory build-up and therefore have advantages in terms of the lean synchronisation of flow (covered in Chapter 12).

OPERATIONS PRINCIPLE

Pull control reduces the build-up on inventory between processes or stages.

Drum, buffer, rope control

The drum, buffer, rope concept comes from the theory of constraints (TOC) described earlier. It is an idea that helps to decide exactly *where* control should occur. Again, the TOC emphasises the role of the bottleneck on work flow. If the bottleneck is the chief constraint, it should be the control point of the whole process. The bottleneck should be the *drum* because it sets the 'beat' for the rest of the process to follow. Because it does not have sufficient capacity, a bottleneck is (or should be) working all the time. Therefore, it is sensible to keep a *buffer* of inventory in front of it to make sure that it always has something to work on. Also, because it constrains the output of the whole process, any time lost at the bottleneck will affect the output from the whole process. So it is not worthwhile for the parts of the process before the bottleneck to work to their full capacity. All they would do is produce work that would accumulate further along in the process up to the point where the bottleneck is constraining flow. Therefore, some form of communication between the bottleneck and the input to the process is needed to make sure that activities before the bottleneck do not overproduce. This is called the *rope* (see Figure 10.12).

OPERATIONS PRINCIPLE

The constraints of bottleneck processes and activities should be a major input to the planning and control activity.

Controlling operations is not always routine

The simple monitoring control model in Figure 10.11 helps us to understand the basic functions of the monitoring and control activity. But it is a simplification. Some simple routine processes may approximate to it, but many other operations do not. In fact, some of the specific criticisms provide a useful set of questions that can be used to assess the degree of difficulty associated with control of any operation:

- Is there consensus over what the operation's objectives should be?
- Are the effects of interventions into the operation predictable?
- Are the operation's activities largely repetitive?

Starting with the first question, are strategic objectives clear and unambiguous? It is not always possible (or necessarily desirable) to articulate every aspect of an operation's objectives in detail. Many operations are just too complex for that. Nor does every senior manager always agree on what the operation's objectives *should* be. Often the lack of a clear objective is because individual managers have different and conflicting interests. In social care organisations, for example, some managers are charged with protecting vulnerable members of society, other with ensuring that public money is not wasted, and yet other may be required to protect the independence of professional staff. At other times objectives are ambiguous because the strategy has to cope with unpredictable changes in the environment making the original objectives redundant. A further assumption in the simplified control model is that there is some reasonable knowledge of how to bring about the desired outcome. That is, when a decision is made, one can predict its effects with a reasonable degree of confidence. In other words, operational control assumes that any interventions that are intended to bring a process back under control will indeed have the intended effect. Yet, this implies that the relationships between the intervention and the resulting consequence within the process are predictable, which in turn assumes that the degree of process knowledge is high. For example, if an organisation decides to relocate in order to be more convenient for its customers, it may or may not prove to be correct. Customers may react in a manner that was not predicted. Even if customers seem initially to respond well to the new location, there may be a lag before negative reactions become evident. In fact many operations decisions are taken about activities about which the cause–effect relationship is only partly understood. The final assumption about control is that control interventions are made in a repetitive way and occur frequently (for example checking on a process, hourly or daily). This means that the operation has the opportunity to learn how its interventions affect the process which considerably facilitates control. However, some control situations are non-repetitive, for example those involving unique services or products. So because the intervention, or the deviation from plan that caused it, may not be repeated, there is little opportunity for learning.

Figure 10.13 illustrates how these questions can form a 'decision tree' type model that indicates how the nature of operations control mat be influenced.[6] Operational control is relatively straightforward: objectives are unambiguous, the effects of interventions are known and activities are repetitive. This type of control can be done using predetermined conventions and rules. There are, however, still some challenges to successful routine control. It needs operational discipline to make sure that control procedures are systematically implemented. The main point though is that any divergence from the conditions necessary for routine control implies a different type of control.

OPERATIONS PRINCIPLE

Planning and control is not always routine, especially when objectives are ambiguous, the effects of interventions into the operation are not predictable and activities not repetitive.

Expert control

If objectives are unambiguous, yet the effects of interventions relatively well understood, but the activity is not repetitive (for example, installing or upgrading software or IT systems) control can be delegated to an 'expert'; someone for whom such activities are repetitive because they have built their knowledge on previous experience elsewhere. Making a success of expert control requires that such experts exist and can be 'acquired' by the firm. It also requires that the expert takes advantage of the control knowledge already present in the firm and integrates his

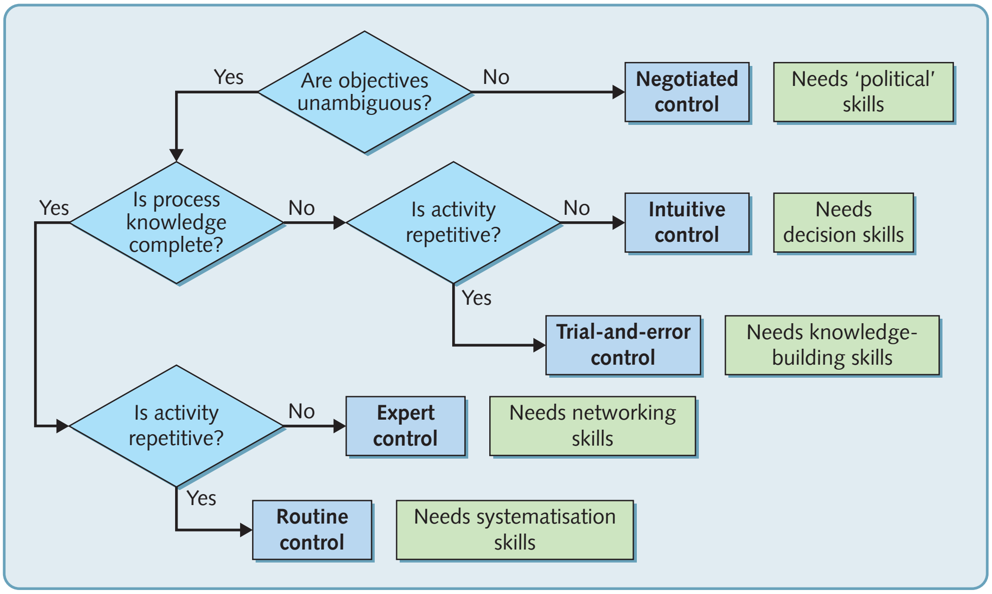

Figure 10.13 Control is not always routine, different circumstances require different types of control

or her 'expert' knowledge with the support that potentially exists internally. Both of these place a stress on the need to 'network', both in terms of acquiring expertise and then integrating that expertise into the organisation.

Trial and error control

If strategic objectives are relatively unambiguous, but the effects of interventions not known, yet the activity is repetitive, the operation can gain knowledge of how to control successfully through its own failures. In other words, although simple prescriptions may not be available in the early stages of making control interventions, the organisation can learn how to do it through experience. For example, if a fast-food chain is opening new stores into new markets, it may not be sure how best to arrange the openings at first. But if the launch is the first of several, the objective must be, not only to make a success of each launch, but equally (or more) important, it must learn from each experience. It is these knowledge-building skills that will ultimately determine the effectiveness of trial and error control.

Intuitive control

If objectives are relatively unambiguous (so it is clear what the operation is trying to do), but the effects of control interventions are not known, and nor are they repetitive, learning by trial and error is not possible. Here control becomes more of an art than a science. And in these circumstances control must be based on the management team using its intuition to make control decisions. Many strategic operations processes fall into this category, for example setting up a strategic supply partnership (see Chapter 7). Objectives are clear (jointly survive in the long term, make and acceptable return, and so on) but, not only are control interventions not repetitive and their effects not fully understood, but sometimes the supplier's interests may be in conflict with yours. Yet, simply stating that 'intuition' is needed in these circumstances is not particularly helpful. Instinct and feelings are, of course, valuable attributes in any management team, but they are the result, at least partly, of understanding how best to organise their shared understanding, knowledge and decision-making skills. It requires thorough decision analysis, not to 'mechanistically' make the decision, but to frame it so that connections can be made, consequences understood and insights gained.

Negotiated control

The most difficult circumstance for strategic control is when objectives are ambiguous. This type of control involves reducing ambiguity in some way by making objectives less uncertain. Sometimes this is done simply by senior managers 'pronouncing' or arbitrarily deciding what objectives *should* be irrespective of opposing views. For example, controlling the activities of a childcare service can involve very different views amongst the professional social workers making day-to-day decisions. Often a negotiated settlement may be sought, which then can become an unambiguous objective. Alternatively outside experts could be used, either to help with the negotiations, or to remove control decisions from those with conflicting views. But, even within the framework of negotiation, there is almost always a political element when ambiguities in objectives exist. Negotiation processes will be, to some extent, dependent on power structures.

Critical commentary

Far from being the magic ingredient which allows operations to fully integrate all their information, ERP is regarded by some as one of the most expensive ways of getting zero or even negative return on investment. For example, the American chemicals giants, Dow Chemical, spent almost half a billion dollars and 7 years implementing an ERP system which became outdated almost as soon as it was implemented. One company, FoxMeyer Drug, claimed that the expense and problems encountered in implementing ERP eventually drove it into bankruptcy. One problem is that ERP implementation is expensive. This is partly because of the need to customise the system, understand its implications on the organisation and train staff to use it. Spending on what some call the ERP ecosystem (consulting, hardware, networking and complementary applications) has been estimated as being twice the spending on the software itself. But it is not only the expense which has disillusioned many companies, it is also the returns they have had for their investment. Some studies show that the vast majority of companies implementing ERP are disappointed with the effect it has had on their businesses. Certainly many companies find that they have to (sometimes fundamentally) change the way they organise their operations in order to fit in with ERP systems. This organisational impact of ERP (which has been described as the corporate equivalent of root-canal work) can have a significantly disruptive effect on the organisation's operations.

● If one accepts only some of the criticisms of ERP, it does pose the question as to why companies have invested such large amounts of money in it. Partly it was the attraction of turning the company's information systems into a 'smooth running and integrated machine'. The prospect of such organisational efficiency is attractive to most managers, even if it does presuppose a very simplistic model of how organisations work in practice. After a while, although organisations could now see the formidable problems in ERP implementation, the investments were justified on the basis that, 'even if we gain no significant advantage by investing in ERP, we will be placed at a disadvantage by *not* investing in it because all our competitors are doing so'. There is probably some truth in this; sometimes businesses have to invest just to stand still.

SUMMARY CHECKLIST

- [ ] Is appropriate effort devoted to planning and controlling the operation's resources and activities?
- [ ] Have any recent failures in planning and control been used to reconsider how the planning and control system operates?
- [ ] Does the system interface with customers to encourage a positive customer experience?
- [ ] Does the planning and control system interface with suppliers to promote a supplier experience that is in your long-term interests?
- [ ] Does the system perform basic planning and control calculations in an appropriate and realistic manner?
- [ ] Is the balance between human and automated decision-making understood and appropriate for the circumstances?
- [ ] How well is resource planning and control information integrated?
- [ ] Have the advantages and disadvantages of moving to a sophisticated (but expensive!) ERP system been investigated?
- [ ] If so, have the possibilities of web integration and supply chain scope been investigated?
- [ ] Are bottlenecks accounted for in the way planning and control decisions are made?
- [ ] If not, have bottlenecks been identified and their effect on the smooth flow of items through the operation been evaluated?

CASE STUDY

subText Studios Singapore

'C.K. One' was clearly upset. Since he had founded *subText* in the fast growing South East Asian computer-generated imaging (CGI) market, three years ago, this was the first time that he had needed to apologise to his clients. In fact, it had been more than an apology; he had agreed to reduce his fee, though he knew that didn't make up for the delay. He admitted that, up to that point, he hadn't fully realised just how much risk there was, both reputational and financial, in failing to meet schedule dates. It wasn't that either he or his team was unaware of the importance of reliability. On the contrary, 'Imagination', 'expertise' and 'reliability', all figured prominently in their promotional literature, mission statements, and so on. It was just that the 'imagination' and 'expertise' parts had seemed to be the things that had been responsible for their success so far. Of course, it had been bad luck that, after more than a year of perfect reliability (not one late job), the two that had been late in the first quarter of 2004 had been particularly critical. *'They were both for new clients'*, said CK, *'And neither of them indicated just how important the agreed delivery date was to them. We should have known, or found out, I admit. But it's always more difficult with new clients, because without a track record with them, you don't really like even to admit the possibility of being late.'*

The company

After studying computer science up to Master's level at the National University of Singapore, C.K. Ong, had worked for four years in CGI workshops in and around the Los Angeles area of California and then taken his MBA at Stanford. It was there that his fellow students had named him CK One, partly because of his fondness for the fragrance and partly because his outgoing leadership talents usually left him as the leader of whatever group he was working in. After that, *'The name just kinda stuck'*, even when he returned to Singapore to start *subText* Studios. While in California, CK had observed that a small but growing part of the market for computer-generated imaging services was in the advertising industry. *'Most CGI work is still connected with the movie industry'*, admitted CK, *'However, two important factors have emerged over the last four or five years. First, the Ad agencies have realised that, with one or two notable exceptions, the majority of their output is visually less arresting than most of the public are used to seeing at the movies. Second, the cost of sophisticated CGI, once something of a barrier to most advertising budgets, is starting to fall rapidly. Partly this is because of cheaper computing power and partly because the scarcity of skilled CGI experts who also have creative talent is starting to rectify itself.'* CK had decided to return to Singapore both for family reasons and because the market in the area was growing quickly and, unlike Hong Kong that had a large movie industry with its ancillary service industries, Singapore had few competitors.

The company was set up on a similar but slightly simpler basis to the companies CK had worked for in California. At the heart of the company were the three 'core' departments that dealt sequentially with each job taken on. These three departments were 'pre-production', 'production' and 'post-production':

- **Pre-production** was concerned with taking and refining the brief as specified by the client, checking with and liaising with the client to iron out any ambiguities in the brief, storyboarding the sequences, and obtaining outline approval of the brief from the client. In addition, pre-production also acted as account liaison with the client and were also responsible for estimating the resources and timing for each job. They also had nominal responsibility for monitoring the job through the remaining two stages, but generally they only did this if the client needed to be consulted during the production and post-production processes. The Supervising

Artists in each department were responsible for the control of the jobs in their departments.

- **Production** involved the creation of the imagery itself. This could be a complex and time-consuming process involving the use of state-of-the-art workstations and CGI software. Around 80 per cent of all production work was carried out in-house, but for some jobs other specialist workshops were contracted. This was only done for work that *subText* either could not do, or would find difficult to do. Contracting was hardly ever used simply to increase capacity because the costs of doing so could drastically reduce margins.
- **Post-production** had two functions: the first was to integrate the visual image sequences produced by production with other effects such as sound effects, music, voice overs, and so on: the second was to cut, edit and generally produce the finished 'product' in the format required by the client.

Each of the three departments employed teams of two people. '*It's a trick I learnt working for a workshop in L.A.*', said CK, '*Two people working together both enhance the creative process (if you get the right two people!), and provide a discipline for each other. Also, it allows for some flexibility both in the mixing of different talents and in making sure that there is always at least one person from the team present at any time who knows the progress and status of any job.*' Pre-production had two teams, production three teams and post-production two teams. In addition, CK himself would sometimes work in the core departments, particularly in pre-production, when he had the time, although this was becoming less common. His main role was in marketing the company's services and business development generally. '*I am the external face of the company and partly my job is to act as a cut-out, particularly for the production and post-production people. The last thing I want is them being disturbed by the clients all the time. I also try and help out around the place when I can, especially with the creative and storyboarding work. The problem with doing this is that, particularly for Pre-production and Post-production, a single extra person assisting, does not always help to get the job done faster. In fact, it can sometimes confuse things and slow things down. This is why, for pre-production and post-production work, one team is always exclusively devoted to one job. We never allow either one team to be working on two jobs at the same time, or have both teams working on one job. It just doesn't work because of the confusion it creates. That doesn't apply to production. Usually (but not always) the production work can be parcelled up so that two or even all three of the teams could be working on different parts of it at the same time. Provided there is close coordination between the teams and provided that they are all committed to pulling it together at the end, there should be a more or less inverse relationship between the number of bodies working on the job and the length of time it takes. In fact, with the infamous 'fifty three slash F' job that's exactly what we had to do. However, notwithstanding what I just said about shortening the time, we probably did lose some efficiency there by having all three teams working on it.*'

'*We pay our teams in the three core department a salary, based on their experience, and a yearly bonus. For that they are expected to, within reason, work until the job is finished. It varies, but most of us work at least ten-hour days relatively frequently. That level of work is factored in to the time estimates we make for each stage of the process. And, although we can be a little inaccurate sometimes, I don't think it's anything to do with a lack of motivation or pace of work. It's just that this type of thing is sometimes difficult to estimate.*'

The fifty three slash F job

The fifty three slash F job, recently finished (late) and delivered to the client (dissatisfied) had been the source of much chaos, confusion and recrimination over the last two or three weeks. Although the job was only three days late, it had caused the client (the Singapore office of a US Advertising agency) to postpone a presentation to its own client. Worse, *subText* had given only five days' notice of late delivery, trying until the last to pull back onto schedule.

The full name of the job that had given them so much trouble was 04/53/F. The 04 related to the year in which the job was started, the 53 was the client's reference number and the F the job identifier (at the start of the year the first job would be labeled A, then B and so forth with AA, BB, etc., being used subsequently). Table 10.1 shows the data for all the jobs started in 2004 up to the current time (day 58; every working day was numbered throughout the year). Figure 10.13 shows the schedule for this period. The job had been accepted on day 18 and had seemed relatively straightforward, although it was always clear that it would be a long production job. It was also clear that time was always going to be tight. There were 32 days in which to finish a job that was estimated to take 30 days.

'*We had been negotiating for this job for two or three weeks and we were delighted to get it. It was important to us because the client, although small in Singapore, had interests all over the world. We saw it as a way into potentially some major business in the future. In hindsight we underestimated how much having three teams working on the production stage of this job at one point or other would increase its complexity. OK, it was not an easy piece of CGI to carry off, but we probably would have been OK if we had organised the CGI stage better. It was also real bad luck that, in our efforts to deliver the fifty three slash F job on time, we also disrupted the fifty*

Table 10.1 *subText* Studios Singapore – Planning data for day 02 to day 58 2004

Job (04)	*Day in*	*Estimated total time*	*Actual total time*	*Due date*	*Actual delivery*	*Pre-prod*		*Prod*		*Post-prod*	
						Est	*Actual*	*Est*	*Actual*	*Est*	*Actual*
06/A	−4	29	30	40	34	6	8	11	10	12	12
11/B	−4	22	24	42	31	4	5.5	7	7.5	11	11
04/C	2	31	30.5	43	40	9	9.5	12	13	10	9
54/D	5	28	34	55	58	10	12	12	17	6	5
31/E	15	34	25	68	57	10	11	12	14	12	-
53/F	18	32	49	50	53	6	10	18	28	8	11
24/G	25	26	20	70	-	9	11	9	9	8	-
22/H	29	32	26	70	-	10	12	14	14	8	-
22/I	33	30	11	75	-	10	11	12	-	8	-
09/J	41	36	14	81	-	12	14	14	-	10	-
20/K	49	40	-	89	-	12	-	14	-	14	-

four slash D job that turned out to be the only other new client we have had this year.' (CK Ong)

The job had proved difficult from the start. The pre-production stage took longer than estimated, mainly because the client's creative team changed just before the start of *subText* beginning the work. But it was the actual CGI itself that proved to be the major problem. Not only was the task intrinsically difficult, it was difficult to parcel it up into separate packages that could be coordinated for working on by the two teams allocated to the job. More seriously, it became apparent within two or three days of starting the production work, that they would need the help of another

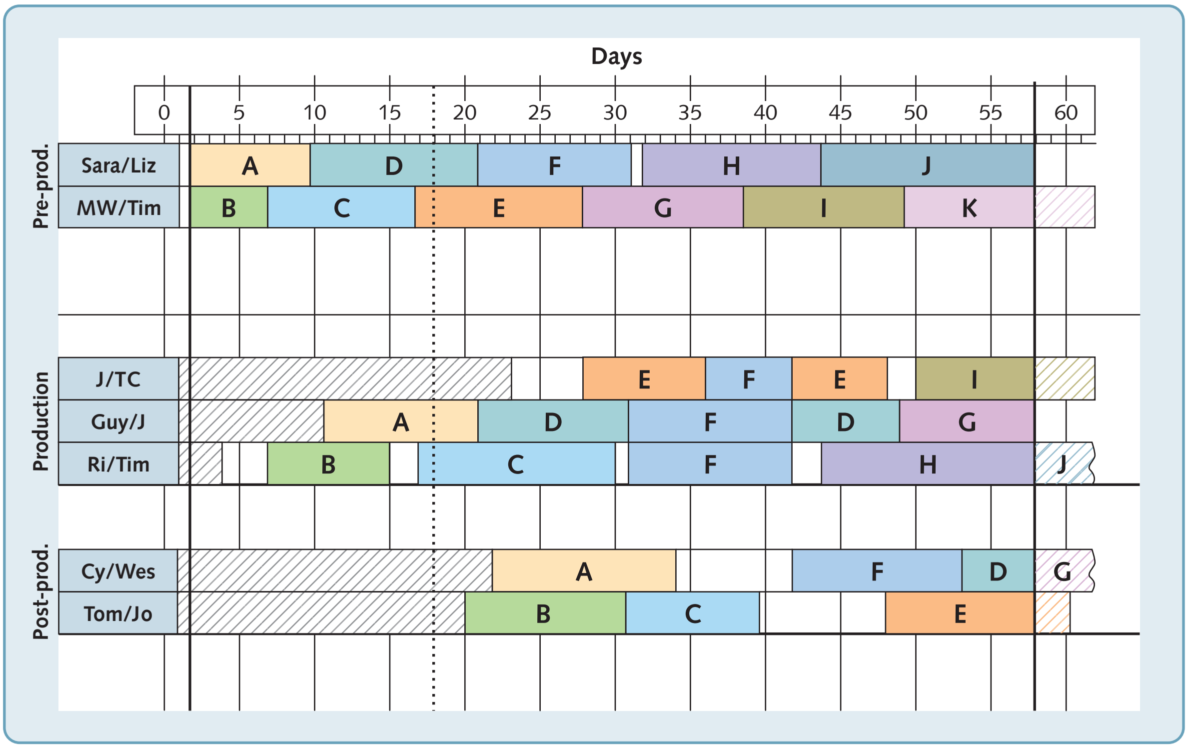

Figure 10.14 *subText* Studios Singapore – Actual schedule for day

studio for some of the effects. Although the other studio was a regular supplier at short notice, this time they were too busy with their own work to help out. Help eventually came from a specialist studio in Hong Kong. *'The subcontracting delay was clearly a problem, but it was only half way through the production phase that we first realised just how much difficulty the fifty three slash F job was in. It was at that stage that we devoted all our production resources to finishing it. Unfortunately, even then, the job was late. The decision eventually to put all three teams on to the fifty three slash F job was not easy because we knew that it would both disrupt other jobs and potentially cause more coordination problems. However, when I accept jobs into production I am accepting that I will try and do whatever it takes to pass it over to Post-production on the date agreed. We did miss out that time, but you can't say we didn't give it everything we've got and technically the job was brilliant, even the client admits that.'* ('TC' Ashwan, Supervising CGI Artist, Production Department)

'No way will be doing that again'

'No way will be doing that again', said CK to the core teams when they met to pick over what had gone wrong. *'We are desperately in need of a more professional approach to keeping track of our activities. There is no point in me telling everyone how good we are if we then let them down. The problem is that I don't want to encourage a 'command and control' culture in the studio. We depend on all staff feeling that they have the freedom to explore seemingly crazy options that may just lead to something real special. We aren't a factory. But we do need to get a grip on our estimating so that we have a better idea of how long each job really will take. After that each of the core departments can be responsible for their own planning.'*

QUESTIONS

1 What went wrong with the fifty three slash F job and how could the company avoid making the same mistakes again?

2 What would you suggest that *subText* do to tighten up their planning and control procedures?

APPLYING THE PRINCIPLES

Some of these exercises can be answered by reading the chapter. Others will require some general knowledge of business activity and some might require an element of investigation. **Model answers for the first two questions are available on the student companion website.** For model answers to the other questions in this section, please ask your tutor/lecturer.

1 Mark Key is an events coordinator for a small company. Returning from his annual holiday in France, he is given six events to plan. He gives them the codes A–F. He needs to decide upon the sequence in which to plan the events and wants to minimise the average time the jobs are tied up in the office and, if possible, meet the deadlines allocated. The six jobs are detailed in Table 10.2.

Table 10.2 The six jobs that Mark has to sequence

Sequence of jobs	*Process time (days)*	*Due date*
A	4	12
B	3	5
C	1	7
D	2	9
E	2	15
F	5	8

(a) Determine a sequence based on using (i) the FIFO rule, (ii) using the Due Date rule and (iii) using the SOT rule.

(b) Which of these sequences gives the most efficient solution and which gives the least lateness?

2 It is week 35 of a busy year at Ashby Architect's and Jo Ashby is facing a big problem. Both her two junior partners have been diagnosed with a serious illness contracted on a trip to scope out a prospective job in Lichtenstein. So Jo has to step in and complete the outstanding jobs that were being worked on by the two juniors. The outstanding jobs are shown in Table 10.3.

Jo has heard that a sequencing rule called the critical ratio (CR) will give efficient results. The priority of jobs using the CR rule is defined by an index computed as follows:

$$CR = \frac{time\ remaining}{work\ days\ remaining} = \frac{due\ date - today's\ date}{work\ days\ remaining}$$

Using this rule, in what priority should Jo give the jobs?

Table 10.3 Outstanding jobs that Jo will need to complete

Job	*Due Date (week)*	*Weeks of work remaining*
Ashthorpe lavatory block	40	2.0
Bubgwitch bus shelters	48	5.0
Crudstone plc HQ	51	3.0
Dredge sewage works	52	8.0

3 It takes six hours for a contract laundry to wash, dry and press (in that order) a batch of overalls. It takes 3 hours to wash the batch, 2 hours to dry it and 1 hour to press it. Usually, each day's batch is collected and ready for processing at 8.00 a.m. and needs to be picked up at 4.00 p.m. The two people who work in the laundry have different approaches to how they schedule the work. One schedules 'forward'. Forward scheduling involves starting work as soon as it arrives. The other person schedules 'backwards'. Backward scheduling involves starting jobs at the last possible moment that will prevent them from being late.

(a) Draw up a schedule indicating the start and finish time for each activity (wash, dry and press) for both forward and backward approaches.

(b) What do you think are the advantages and disadvantages of these two approaches?

4 Read the following descriptions of two cinemas.
Kinepolis in Brussels is one the largest cinema complexes in the world, with 28 screens, a total of 8,000 seats, and four showings of each film every day. It is equipped with the latest projection technology. All the film performances are scheduled to start at the same times every day: 4 p.m., 6 p.m., 8 p.m. and 10.30 p.m. Most customers arrive during the 30 minutes before the start of the film. Each of the 18 ticket desks has a networked terminal and a ticket printer. For each customer, a screen code is entered to identify and confirm seat availability of the requested film. Then the number of seats required is entered, and the tickets are printed, though these do not allocate specific seat positions. The operator then takes payment by cash or credit card and issues the tickets. This takes an average of 19.5 seconds, a further 5 seconds is needed for the next customer to move forward. An average transaction involves the sale of approximately 1.7 tickets.

The UCI cinema in Birmingham has eight screens. The cinema incorporates many 'state-of-the-art' features, including the high-quality THX sound system, fully computerised ticketing and a video games arcade off the main hall. In total the eight screens can seat 1,840 people; the capacity (seating) of each screen varies, so the cinema management can allocate the more popular films to the larger screens and use the smaller screens for the less popular films. The starting times of the eight films at UCI are usually staggered by 10 minutes, with the most popular film in each category (children's, drama, comedy, etc.) being scheduled to run first. Because the films are of different durations, and since the manager must try to maximise the utilisation of the seating, the scheduling task is complex. Ticket staff are continually aware of the remaining capacity of each 'screen' through their terminals. There are up to four ticket desks open at any one time. The target time per overall transaction is 20 seconds. The average number of ticket sales per transaction is 1.8. All tickets indicate specific seat positions, and these are allocated on a first-come-first-served basis.

(a) What are the main differences between the two cinemas from the perspectives of their operations managers?

(b) What are the advantages and disadvantages of the two different methods of scheduling the films onto the screens?

(c) Find out the running times and classification of eight popular films. Try to schedule these onto the UCI Solihull screens, taking account of what popularity you might expect at different times. Allow 20 minutes for emptying, cleaning and admitting the next audience, and 15 minutes for advertising, before the start of the film.

5 Think through the following three brief examples. What type of control (according to Figure 10.13) do you think they warrant?

1 The Games Delivery Authority (GDA) was a public body responsible for developing and building the new venues and infrastructure for the 'International Games' and their use after the event. The GDA appointed a consortium responsible for the overall programme's quality, delivery and cost in addition to health and safety, sustainability, equality and diversity targets. The Games Park was a large construction programme, spreading across five separate local government areas, including transport developments, retail areas and local regeneration projects. Sustainability was central to the development. '*Sustainability' was ingrained into our thinking – from the way we planned, built and worked, to the way we play, socialise and travel.'* To ensure they stuck to commitments, the GDA set up an independent body to monitor the project. All potential contractors tendering for parts of the project were aware that a major underlying objective of the Games initiative was regeneration. The Games site was to be built on highly industrialised and contaminated land.

2 The supermarket's new logistics boss was blunt in his assessment of its radical supply chain implementation. '*Our rivals have watched in utter disbelief*', he said. '*Competitors looked on in amazement as we poured millions into implementing new IT systems and replaced 21 depots with a handful of giant automated 'fulfilment factories*'. '*In hindsight, the heavy reliance on automation was a big mistake, especially for fast moving goods*', said the company's CIO. '*When a conventional facility goes wrong, you have lots of options. You have flexibility to deal with issues. When an automated 'fulfilment factory' goes wrong, frankly, you're stuck*.' Most damning was the way that the supermarket pressed on with the implementation of the automated facilities before proving that the concept worked at the first major site. '*I'd have at least proved that one of them worked before building the other three*', he said. '*Basically, the whole company was committed to doing too much, too fast, trying to implement a seven-year strategy in a three-year timescale*.'

3 '*It's impossible to overemphasise just how important this launch is to our future*', said the CEO. '*We have been losing market share for seven quarters straight. However, we have very high hopes for the new XC10 unit*.' And most of the firm's top management team agreed with her. Clearly the market had been maturing for some time now, and was undoubtedly getting more difficult. New product launches from competitors, had been eroding both market share. Yet competitors' products, at best, simply matched the firm's offerings in all benchmark tests. '*Unless someone comes up with a totally new technology, which is very unlikely, it will be a matter of making marginal improvements in product performance and combing this with well targeted and coordinated marketing. Fortunately, we are good at both of these. We know this technology, and we know these markets. We are also clear what role the new XC10 should play. It needs to consolidate our market position as the leader in this field, half the slide in market share, and re-establish our customers' faith in us. Margins, at least in the short term, are less important*.'

Notes on chapter

1. Sources include: Masresha, A. K. (2011) *Aircraft Scheduling,* Lambert Academic Publishing; Farman, J. (1999) 'Les Coulisses du Vol', Air France. Talk presented by Richard E. Stone, NorthWest Airlines at the IMA Industrial Problems Seminar, 1998.
2. Source: Interview with Joanne Cheung, Steve Deeley and other staff at Godfrey Hall, BMW Dealership, Coventry.
3. Kanaracus, C. (2008) 'Waste Management sues SAP over ERP implementation', *InfoWorld,* 27 March.
4. Coldrick A, Ling, D. and Turner, C. (2003) 'Evolution of sales and operations planning', StrataBridge Working Paper.
5. Goldratt, E.Y. and Cox, J. (1984) *The Goal,* North River Press.
6. Based on an original model described in Hofstede, G. (198l) 'Management control of public and not-for-profit activities', *Accounting, Organizations and Society,* vol. 6, no. 3, 193–211.

TAKING IT FURTHER

Akhtar, J. (2013) ***Production Planning and Control with SAP ERP, SAP Press,*** *Comprehensive, but specialised and focused on one particular ERP vendor (SAP).*

Chapman, S.N. (2005) ***Fundamentals of Production Planning and Control, Pearson.*** *A good basic textbook.*

Dickersbach, J. and Keller, G. (2010) ***Production Planning and Control with SAP ERP (2nd edn), Sap Press/Galileo Press.*** *Technical but thorough.*

Goldratt, E.Y. and Cox, J. (2004) ***The Goal: A Process of Ongoing Improvement, Gower Publishing Ltd.*** *Don't read this if you like good novels but do read this if you want an enjoyable way of understanding some of the complexities of scheduling. It particularly applies to the drum, buffer, rope concept described in this chapter and it also sets the scene for the discussion of OPT.*

Magal, S.R. and Word, J. (2010) ***Integrated Business Processes with ERP Systems, John Wiley & Sons.*** *It's written in partnership with SAP, the leading seller of ERP systems, but it does covers all of the key processes supported by modern ERP systems.*

Vollmann, T.E., Berry, W.L. and Whybark, D.C. (2010) ***Manufacturing Planning and Control System for Supply Chain Management (6th edn), McGraw-Hill Higher Education.*** *This is the bible of production planning and control. It deals with all the issues in this part of this textbook.*

SUPPLEMENT TO CHAPTER 10

Materials requirements planning (MRP)

Materials requirements planning (MRP) is an approach to calculating how many parts or materials of particular types are required and what times they are required. This requires data files which, when the MRP program is run, can be checked and updated. Figure 10.15 shows how these files relate to each other. The first inputs to materials requirements planning are customer orders and forecast demand. MRP performs its calculations based on the combination of these two parts of future demand. All other requirements are derived from, and dependent on, this demand information.

Master production schedule

The master production schedule (MPS) forms the main input to materials requirements planning and contains a statement of the volume and timing of the end products to be made. It drives all the production and supply activities that eventually will come together to form the end products. It is the basis for the planning and utilisation of labour and equipment, and it determines the provisioning of materials and cash. The MPS should include all sources of demand, such as spare parts, internal production promises, and so on. For example, if a manufacturer of earth excavators plans an exhibition of its products and allows a project team to raid the stores so that it can build two pristine examples to be exhibited, this is likely to leave the factory short of parts. MPS can also be used in service organisations. For example, in a hospital theatre there is a master schedule that contains a statement of which operations are planned and when. This can be used to provision materials for the operations, such as the sterile instruments, blood and dressings. It may also govern the scheduling of staff for operations.

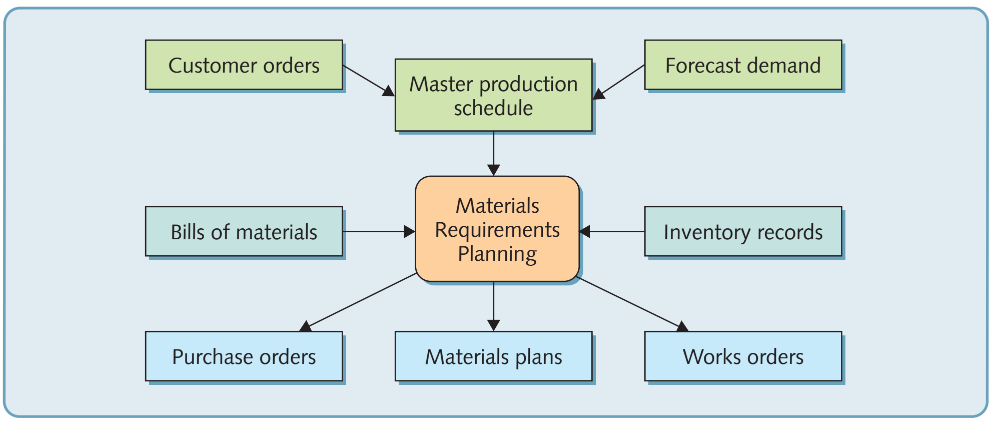

Figure 10.15 **Materials requirements planning (MRP) schematic**

The master production schedule record

Master production schedules are time-phased records of each end product, which contain a statement of demand and currently available stock of each finished item. Using this information, the available inventory is projected ahead in time. When there is insufficient inventory to satisfy forward demand, order quantities are entered on the master schedule line. Table 10.4 is a simplified example of part of a master production schedule for one item. In the first row the known sales orders and any forecast are combined to form 'Demand'. The second row, 'Available', shows how much inventory of this item is expected to be in stock at the end of each weekly period. The opening inventory balance, 'On-hand', is shown separately at the bottom of the record. The third row is the master production schedule, or MPS; this shows how many finished items need to be completed and available in each week to satisfy demand.

Chase or level master production schedules

In the example in Table 10.4, the MPS increases as demand increases and aims to keep available inventory at 0. The master production schedule is 'chasing' demand (see Chapter 10) and so adjusting the provision of resources. An alternative 'levelled' MPS for this situation is shown in Table 10.5. Level scheduling involves averaging the amount required to be completed to smooth out peaks and troughs; it generates more inventory than the previous MPS.

Available to promise (ATP)

The master production schedule provides the information to the sales function on what can be promised to customers and when delivery can be promised. The sales function can load known sales orders against the master production schedule and keep track of what is

Table 10.4 Example of a master production schedule

		Week number								
		1	*2*	*3*	*4*	*5*	*6*	*7*	*8*	*9*
Demand		10	10	10	10	15	15	15	20	20
Available		20	10	0	0	0	0	0	0	0
MPS		0	0	10	10	15	15	15	20	20
On-hand	30									

Table 10.5 Example of a 'level' master production schedule

		Week number								
		1	*2*	*3*	*4*	*5*	*6*	*7*	*8*	*9*
Demand		10	10	10	10	15	15	15	20	20
Available		31	32	33	34	30	26	22	13	4
MPS		11	11	11	11	11	11	11	11	11
On-hand	30									

Table 10.6 Example of a level master production schedule including available to promise

		Week number								
		1	*2*	*3*	*4*	*5*	*6*	*7*	*8*	*9*
Demand		10	10	10	10	15	15	15	20	20
Sales orders		10	10	10	8	4				
Available		31	32	33	34	30	26	22	13	4
ATP		31	1	1	3	7	11	11	11	11
MPS		11	11	11	11	11	11	11	11	11
On-hand	30									

available to promise (ATP) (see Table 10.6). The ATP line in the master production schedule shows the maximum that is still available in any one week, against which sales orders can be loaded.

The bill of materials (BOM)

From the master schedule, MRP calculates the required volume and timing of assemblies, sub-assemblies and materials. To do this, it needs information on what parts are required for each product. This is called the 'bill of materials'. Initially it is simplest to think about these as a product structure. The product structure in Figure 10.16 is a simplified structure showing the parts required to make a simple board game. Different 'levels of assembly' are

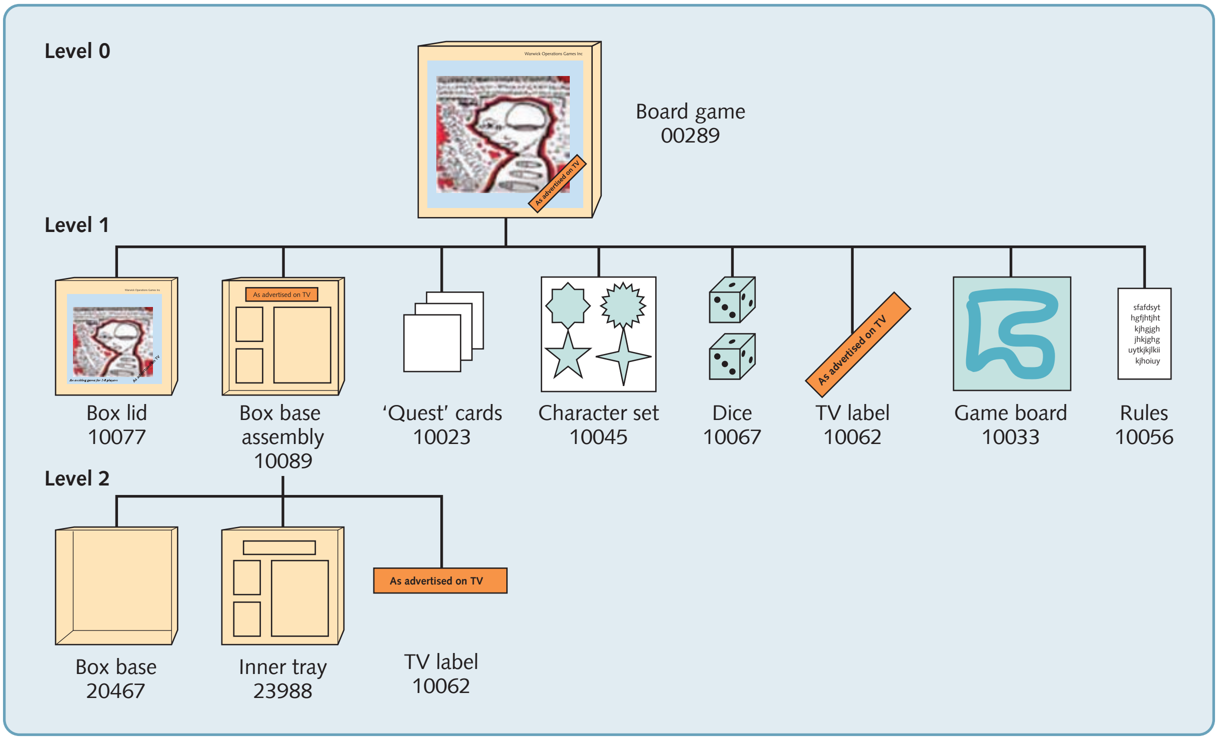

Figure 10.16 Product structure for a simple board game

Table 10.7 Indented bill of materials for board game

Part number: 00289

Description: Board game

Level: 0

Level	*Part number*	*Description*	*Quantity*
0	00289	Board game	1
.1	10077	Box lid	1
.1	10089	Box base assy	1
..2	20467	Box base	1
..2	10062	TV label	1
..2	23988	Inner tray	1
.1	10023	Quest cards set	1
.1	10045	Character set	1
.1	10067	Die	2
.1	10062	TV label	1
.1	10033	Game board	1
.1	10056	Rules booklet	1

shown with the finished product (the boxed game) at level 0, the parts and sub-assemblies that go into the boxed game at level 1, the parts that go into the sub-assemblies at level 2, and so on.

A more convenient form of the product structure is the 'indented bill of materials'. Table 10.7 shows the whole indented bill of materials for the board game. The term 'indented' refers to the indentation of the level of assembly, shown in the left-hand column. Multiples of some parts are required; this means that MRP has to know the required number of each part to be able to multiply up the requirements. Also, the same part (for example, the TV label, part number 10062) may be used in different parts of the product structure. This means that MRP has to cope with this commonality of parts and, at some stage, aggregate the requirements to check how many labels in total are required.

Inventory records

MRP calculations need to recognise that some required items may already be in stock. So, it is necessary, starting at level 0 of each bill, to check how much inventory is available of each finished product, sub-assembly and component, and then to calculate what is termed the 'net' requirements, that is the extra requirements needed to supplement the inventory so that demand can be met. This requires that three main inventory records are kept: the item master file, which contains the unique standard identification code for each part or component; the transaction file, which keeps a record of receipts into stock, issues from stock and a running balance; and the location file, which identifies where inventory is located.

The MRP netting process

The information needs of MRP are important, but it is not the 'heart' of the MRP procedure. At its core, MRP is a systematic process of taking this planning information and calculating the volume and timing requirements that will satisfy demand. The most important element of this is the MRP netting process.

Figure 10.17 illustrates the process that MRP performs to calculate the volumes of materials required. The master production schedule is 'exploded', examining the implications of the schedule through the bill of materials, checking how many sub-assemblies and parts are required. Before moving down the bill of materials to the next level, MRP checks how many of the required parts are already available in stock. It then generates 'works orders', or requests, for the net requirements of items. These form the schedule, which is again exploded through the bill of materials at the next level down. This process continues until the bottom level of the bill of materials is reached.

Back-scheduling

In addition to calculating the volume of materials required, MRP also considers when each of these parts is required, that is, the timing and scheduling of materials. It does this by a process called back-scheduling, which takes into account the lead time (the time allowed for completion of each stage of the process) at every level of assembly. Again using the example of the board game, assume that 10 board games are required to be finished by a notional planning day which

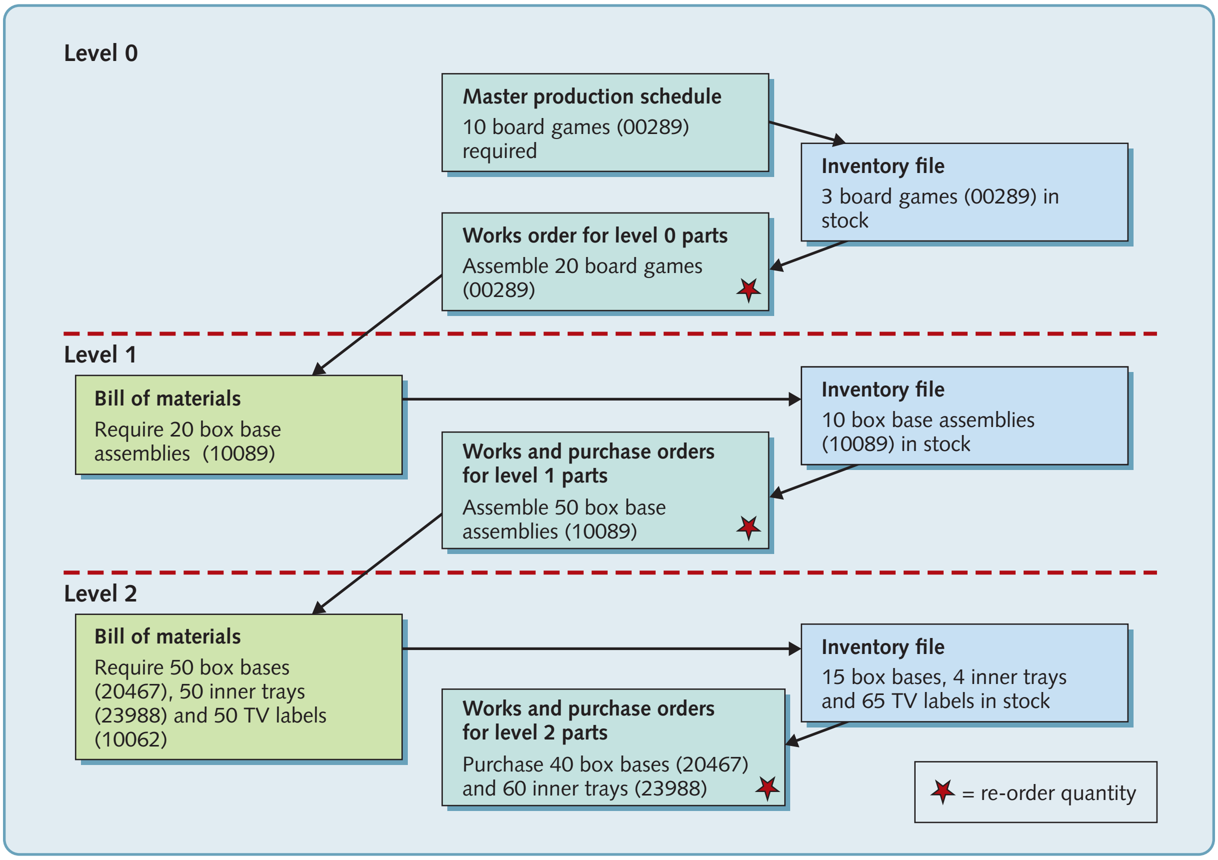

Figure 10.17 The MRP netting calculations for the simple board game

Table 10.8 Back-scheduling of requirements in MRP

Part no.	*Description*	*Inventory on-hand day 0*	*Lead time (days)*	*Re-order quantity*
00289	Board game	3	2	20
10077	Box lid	4	8	25
10089	Box base assy	10	4	50
20467	Box base	15	12	40
23988	Inner tray	4	14	60
10062	TV label	65	8	100
10023	Quest cards set	4	3	50
10045	Character set	46	3	50
10067	Die	22	5	80
10033	Game board	8	15	50
10056	Rules booklet	0	3	80

we will term day 20. To determine when we need to start work on all the parts that make up the game, we need to know all the lead times that are stored in MRP files for each part (see Table 10.8).

Using the lead-time information, the programme is worked backwards to determine the tasks that have to be performed and the purchase orders that have to be placed. Given the lead times and inventory levels shown in Table 10.6, the MRP records shown in Figure 10.18 can be derived.

MRP capacity checks

The MRP process needs a feedback loop to check whether a plan was achievable and whether it has actually been achieved. Closing this planning loop in MRP systems involves checking production plans against available capacity and, if the proposed plans are not achievable at any level, revising them All but the simplest MRP systems are now closed-loop systems. They use three planning routines to check production plans against the operation's resources at three levels:

1. Resource requirements plans (RRPs) - involve looking forward in the long term to predict the requirements for large structural parts of the operation, such as the numbers, locations and sizes of new plants.
2. Rough-cut capacity plans (RCCPs) – are used in the medium to short term, to check the master production schedules against known capacity bottlenecks, in case capacity constraints are broken. The feedback loop at this level checks the MPS and key resources only.
3. Capacity requirements plans (CRPs) – look at the day-to-day effect of the works orders issued from the MRP on the loading individual process stages

00289: Treasure Hunt game Assembly lead time = 2 Re-order quantity = 20

Day Number:	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Requirements Gross																					10
Scheduled Receipts																					
On hand Inventory	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	13
Planned Order Release																			20		

10077: Box lid Purchase lead time = 8 Re-order quantity = 25

Day Number:	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Requirements Gross																			20		
Scheduled Receipts																					
On hand Inventory	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4	9	9	9
Planned Order Release											25										

10089: Box base assembly Assembly lead time = 4 Re-order quantity = 50

Day Number:	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Requirements Gross																			20		
Scheduled Receipts																					
On hand Inventory	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	10	40	40	40
Planned Order Release															50						

20467: Box base Purchase lead time = 12 Re-order quantity = 40

Day Number:	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Requirements Gross															50						
Scheduled Receipts																					
On hand Inventory	15	15	15	15	15	15	15	15	15	15	15	15	15	15	5	5	5	5	5	5	5
Planned Order Release			40																		

23988: Inner tray Purchase lead time = 14 Re-order quantity = 60

Day Number:	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Requirements Gross															50						
Scheduled Receipts																					
On hand Inventory	4	4	4	4	4	4	4	4	4	4	4	4	4	4	14	14	14	14	14	14	14
Planned Order Release	60																				

10062: TV label Purchase lead time = 8 Re-order quantity = 100

Day Number:	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Requirements Gross															50				20		
Scheduled Receipts																					
On hand Inventory	65	65	65	65	65	65	65	65	65	65	65	65	65	65	15	15	15	15	95	95	95
Planned Order Release											100										

10023: Quest card set Purchase lead time = 3 Re-order quantity = 50

Day Number:	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Requirements Gross																			20		
Scheduled Receipts																					
On hand Inventory	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4	4	34	34	34
Planned Order Release																50					

10045: Character set Purchase lead time = 3 Re-order quantity = 50

Day Number:	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Requirements Gross																			20		
Scheduled Receipts																					
On hand Inventory	46	46	46	46	46	46	46	46	46	46	46	46	46	46	46	46	46	46	26	26	26
Planned Order Release																					

10067: Dice Purchase lead time = 5 Re-order quantity = 80

Day Number:	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Requirements Gross																			40		
Scheduled Receipts																					
On hand Inventory	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3	13
Planned Order Release														80							

10033: Game board Purchase lead time = 15 Re-order quantity = 50

Day Number:	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Requirements Gross																			20		
Scheduled Receipts																					
On hand Inventory	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	8	38	38	38
Planned Order Release				50																	

10056: Rules booklet Purchase lead time = 3 Re-order quantity = 80

Day Number:	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Requirements Gross																			20		
Scheduled Receipts																					
On hand Inventory	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	60	60	60
Planned Order Release																80					

Figure 10.18 Extract from the MRP records for the simple board game (lead times indicated by arrows > <)

11 Lean synchronisation

Introduction

Lean synchronisation has the aim of using 'lean' or 'just-in-time' (JIT) principles to supply perfect quality products and services in synchronisation with the demand for them, with zero waste, and at low cost. Once a radical departure from traditional operations practice, these principles have transformed the way leading organisations think about their operations. Although the topic is sometime narrowly treated as a manufacturing phenomenon (perhaps unsurprisingly given the pioneering role of Toyota in lean management), lean synchronisation principles can be applied, should be applied and are applied across all sectors, including finance, healthcare, IT, retailing, construction, agriculture and the public sector. Figure 11.1 shows the position of the ideas described in this chapter in the general model of operations management.

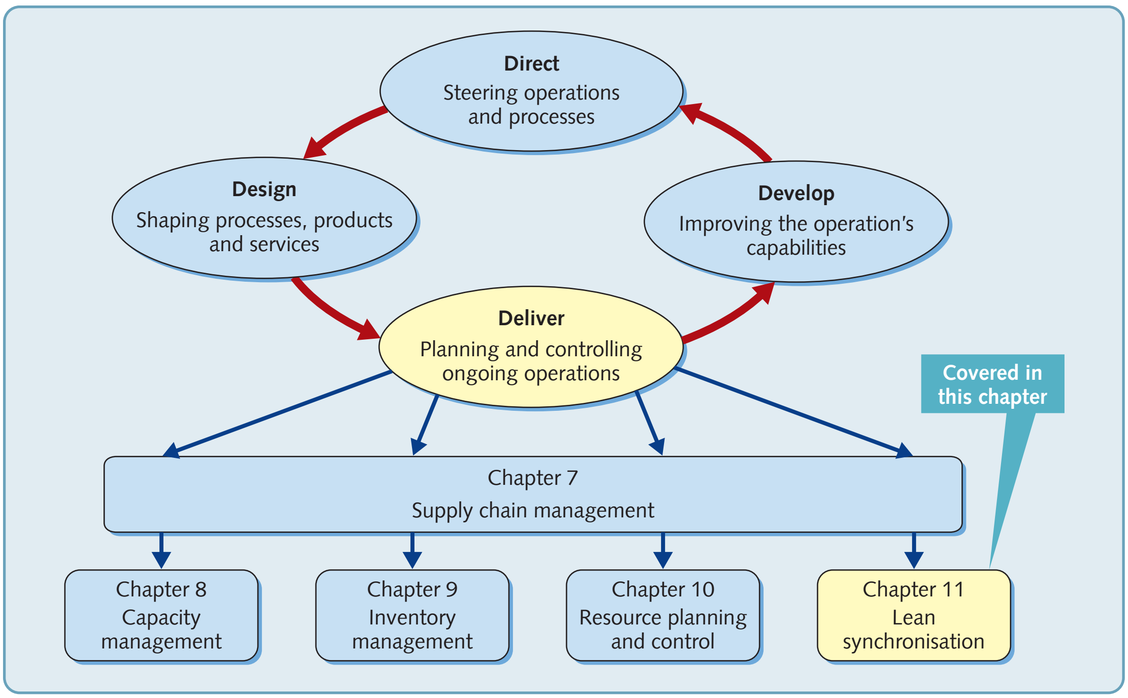

Figure 11.1 Lean synchronisation has the aim of achieving a flow of products and services that always delivers exactly what customers want, in exact quantities, exactly when needed, exactly where required and at the lowest possible cost

EXECUTIVE SUMMARY

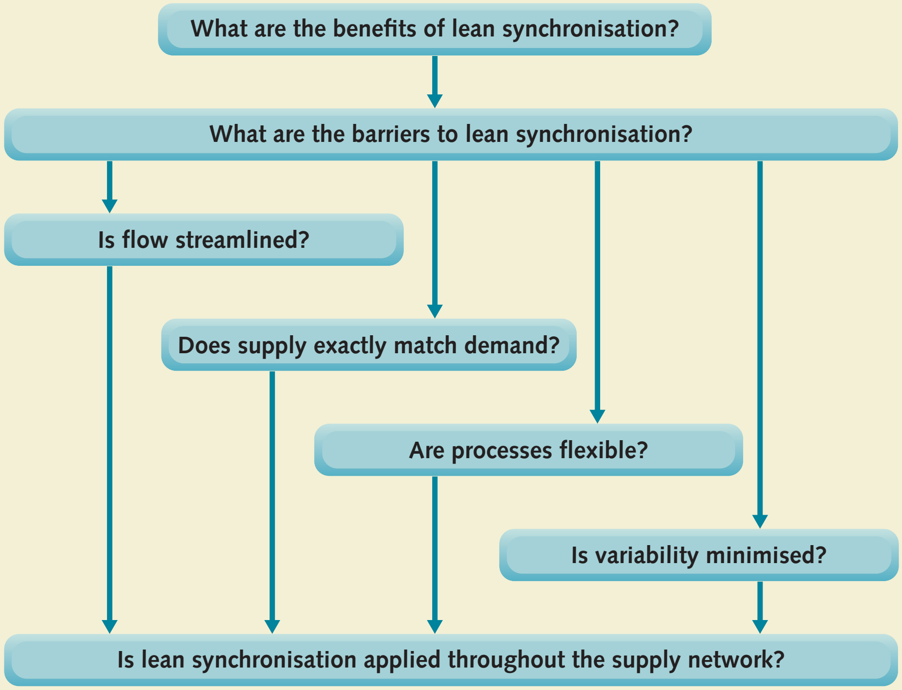

What are the benefits of lean synchronisation?

The focus of lean synchronisation is to achieve a flow of products and services that delivers exactly what customers want, in exact quantities, exactly when needed, exactly where required and at the lowest possible cost. It is a concept that is almost synonymous with terms such as 'just-in-time' (JIT) and 'lean operations principles'. The central idea is that if products, customers, or information flow smoothly, not only is throughput time reduced, but the negative effects of in-process inventory or queues are avoided. Inventories and queues are seen as obscuring the problems that exist within processes and therefore inhibiting process improvement.

What are the barriers to lean synchronisation?

The aim of lean synchronisation can be inhibited in three ways. First is the failure to eliminate waste in all parts of the operation; and the causes of waste are more extensive than is generally understood. The second is a failure to involve all the people within the operation in the shared task of smoothing flow and eliminating waste. Japanese proponents of lean synchronisation often use a set of 'basic working practices' to ensure involvement. Third is the failure to adopt continuous improvement principles. Because pure lean synchronisation is an aim rather than something that can be implemented quickly, it requires the continual application of incremental improvement steps to reach its potential.

Is flow streamlined?

Long process routes are wasteful and cause delay and inventory build-up. Physically reconfiguring processes to reduce distance travelled and aid cooperation between staff can help to streamline flow. Similarly, ensuring flow visibility helps to make improvement to flow easier. Sometimes this can involve small-scale technologies that can reduce fluctuations in flow volume.

Does supply exactly match demand?

The aim of lean synchronisation is to meet demand exactly; neither too much nor too little, and only when it is needed. Pull control principles are typically used to achieve this goal. The most common method of doing this is the use of kanbans - simple signalling devices that prevent the accumulation of excess inventory.

Are processes flexible?

Responding exactly to demand only when it is needed often requires a degree of flexibility in processes, both to cope with unexpected demand and to allow processes to change between different activities without excessive delay. This often means reducing changeover times in technologies.

Is variability minimised?

Variability in processes disrupts flow and prevents lean synchronisation. Variability includes quality variability and schedule variability. Statistical process control (SPC) principles are useful in reducing quality variability. The use of levelled scheduling and mixed modelling can be used to reduce flow variability and total productive maintenance (TPM) can reduce the variability caused by breakdowns.

Is lean synchronisation applied throughout the supply network?

The same benefits and principles of lean synchronisation that apply within operations can also apply across supply networks. This is more difficult, partly because of the complexity of flow and partly because supply networks are prone to the type of unexpected fluctuations that are easier to control within operations.

DIAGNOSTIC QUESTION

What are the benefits of lean synchronisation?

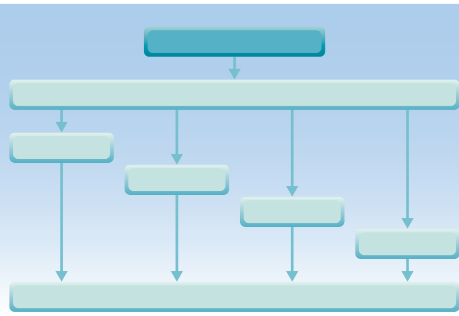

Synchronisation means that the flow of products and services always delivers exactly what customers want (perfect quality), in exact quantities (neither too much nor too little), exactly when needed (not too early or too late) and exactly where required (not to the wrong location). *Lean* synchronisation is to do all this at the lowest possible cost. It results in customers, products and information flowing rapidly and smoothly through processes, operations and supply networks. As such, it is not an idea that is relevant only to one context. While many of the principles of lean described here emerged in manufacturing organisations initially, they are nearly all applicable to other sectors. In this chapter, we will provide examples of organisations in a wide range of sectors that apply facets of lean to their operations and supply networks. Many of the examples of lean philosophy and lean techniques in service industries are directly analogous to those found in manufacturing, because physical items are being moved or processed in some way. For example, supermarkets usually replenish their shelves only when customers have taken sufficient products off the shelf. The movement of goods from the 'back office' store to the shelf is triggered only by the 'empty-shelf' demand signal. Construction companies increasingly make a rule of only calling for material deliveries to its sites the day before materials are needed. This reduces clutter and the chances of theft. Both are examples of the application of *pull control* principles. Other examples of lean concepts and methods apply even when most of the service elements are intangible. For example, new publishing technologies allow professors to assemble printed and e-learning course material customised to the needs of individual courses or even individual students. Here, we see the lean principles of flexibility and small batch sizes allowing customisation and rapid delivery.

The benefits of synchronised flow

The best way to understand how lean synchronisation differs from more traditional approaches to managing flow is to contrast the two simple processes in Figure 11.2. The traditional approach assumes that each stage in the process will place its output in an inventory or queue that 'buffers' that stage from the next one downstream in the process. The next stage down will then (eventually) take outputs from this buffer, process them and pass them through to the next buffer inventory or queue. These buffers are there to 'insulate' each stage from its neighbours, making each stage relatively independent so that if, for example, stage A stops operating for some reason, stage B can continue, at least for a time. The larger is the buffer between stages in the process, the greater the degree of insulation between these stages. This insulation has to be paid for in terms of inventory or queues and slow throughput times because products, customers or information will spend time waiting between stages in the process.

But, the main 'learning' argument against this traditional approach lies in the very conditions it seeks to promote, namely the insulation of the stages from one another. When a problem occurs at one stage, the problem will not immediately be apparent elsewhere in the system. The responsibility for solving the problem will be centred largely on the people within that stage, and the consequences of the problem will be prevented from spreading to the whole system.

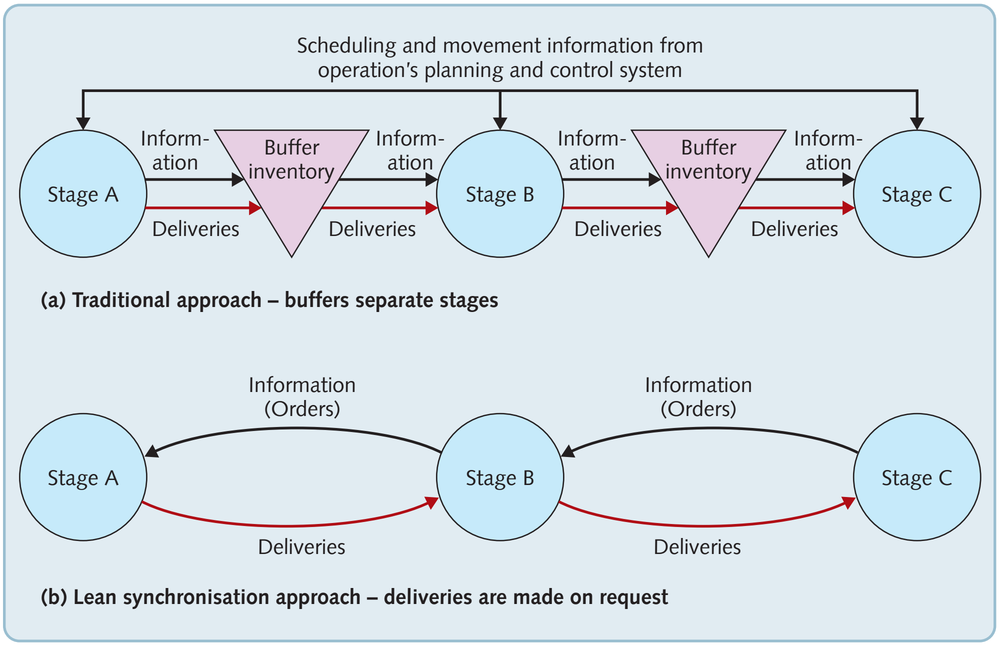

Figure 11.2 **(a) Traditional and (b) lean synchronised flow between stages**

However, contrast this with the pure lean synchronised process illustrated in Figure 11.2. Here products, customers, or information are processed and then passed directly to the next stage 'just-in-time' for them to be processed further. Problems at any stage have a very different effect in such a system. Now if stage A stops processing, stage B will notice immediately and stage C very soon after. Stage A's problem is now quickly exposed to the whole process, which is immediately affected by the problem. This means that the responsibility for solving the problem is no longer confined to the staff at stage A. Everyone now shares the problem, considerably improving the chances of the problem being solved, if only because it is now too important to be ignored. In other words, by preventing products, customers, or information resources accumulating between stages, the operation has increased the chances of the intrinsic efficiency of the operation being improved.

Non-synchronised approaches seek to encourage efficiency by protecting each part of the process from disruption. The lean synchronised approach takes the opposite view. Exposure of the system (although not suddenly, as in our simplified example) to problems can both make them more evident and change the 'motivation structure' of the whole system towards solving the problems. Lean synchronisation sees accumulations of inventories, be they product, customer, of information inventories, as a 'blanket of obscurity' that lies over the system and prevents problems being noticed. Table 11.1 shows the operational consequences of the build-up of these different inventory types.

OPERATIONS PRINCIPLE

Buffer inventories (product, customer, or information) used to insulate stages localise the motivation to improve.

The river and rocks analogy

The idea of obscuring the effects of inventory is often illustrated diagrammatically, as in Figure 11.3. The many problems of the operation are shown as rocks in a riverbed that cannot be seen because of the depth of the water. The water in this analogy represents the inventory in the operation. Yet, even though the rocks cannot be seen, they slow the progress of the river's flow and cause turbulence. Gradually reducing the depth of the water (inventory) exposes the worst of the problems which can be resolved, after which the water is lowered further, exposing more problems, and so on. The same argument applies for flow between whole

Table 11.1 Operational implications of different inventory types

	Inventory		
	Material (queue of material)	*Information (queue of information)*	*Customers (queue of people)*
Cost	Ties up working capital	Less current information and so worth less	Wastes customer's time
Space	Needs storage space	Needs memory capacity	Need waiting area
Quality	Defects hidden, possible damage	Defects hidden, possible data corruption	Gives negative perception
Decoupling	Makes stages independent	Makes stages independent	Promotes job specialisation / fragmentation
Utilisation	Stages kept busy by work-in-progress	Stages kept busy by work in data queues	Servers kept busy by waiting customers
Coordination	Avoids need for synchronisation	Avoids need for straight through processing	Avoids having to match supply and demand

Source: Adapted from Fitzsimmons, J.A. (1990) 'Making Continual Improvement: A Competitive Strategy for Service Firms' *in* Bowen, D.E., Chase, R.B., Cummings, T.G. and Associates (eds) *Service Management Effectiveness*, Jossey-Bass.

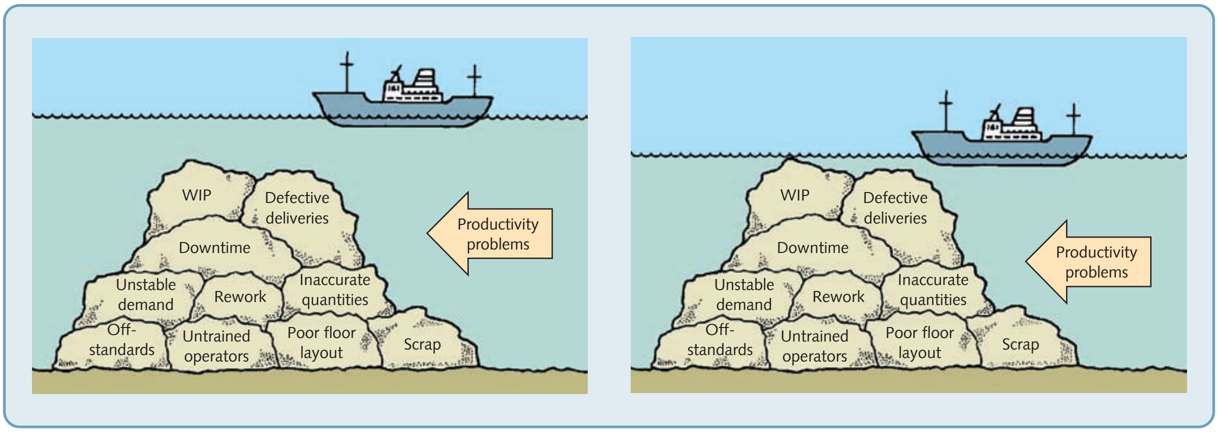

Figure 11.3 Reducing the level of inventory or queues (water) allows operations management (the ship) to see the problems in the operation (the rocks) and work to reduce them

processes or whole operations. For example, stages A, B and C in Figure 11.2 could be a supplier operation, a manufacturer and a customer's operation, respectively.

Synchronisation, 'lean' and 'just-in-time'

Different terms are used to describe what here we call lean synchronisation. Our shortened definition – *'lean synchronisation aims to meet demand instantaneously, with perfect quality and no waste'* – could also be used to describe the general concept of 'lean', or 'just-in-time' (JIT). The concept of 'lean' stresses the elimination of waste, while 'just-in-time' emphasises the idea of producing items only when they are needed. But all three concepts overlap to a large degree, and no definition fully conveys the full implications for operations practice. Here we use the term lean synchronisation because it best describes the impact of these ideas on flow and delivery.

Two operations that have implemented lean synchronisation are briefly described below. One is the company that is generally credited with doing most to develop the whole concept; the other is a not-for-profit hospital that, nevertheless, has derived benefits from adopting some of the principles.

EXAMPLE

Toyota[1]

Seen as the leading practitioner and the main originator of the lean approach, the Toyota Motor Company has progressively synchronised all its processes simultaneously to give high-quality, fast throughput and exceptional productivity. It has done this by developing a set of practices that has largely shaped what we now call 'lean' or 'just-in-time' but which Toyota calls the Toyota production system (TPS). The TPS has two themes, 'just-in-time' and 'jidoka'. Just-in-time is defined as the rapid and coordinated movement of parts throughout the production system and supply network to meet customer demand. It is operationalised by means of *heijunka* (levelling and smoothing the flow of items), *kanban* (signalling to the preceding process that more parts are needed) and *nagare* (laying out processes to achieve smoother flow of parts throughout the production process). *Jidoka* is described as 'humanising the interface between operator and machine'. Toyota's philosophy is that the machine is there to serve the operator's purpose. The operator should be left free to exercise their judgement. Jidoka is operationalised by means of a fail-safe (or machine jidoka), line-stop authority (or human jidoka) and visual control (at-a-glance status of production processes and visibility of process standards).

Toyota believe that both just-in-time and jidoka should be applied ruthlessly to the elimination of waste, where waste is defined as 'anything other than the minimum amount of equipment, items, part and workers that are absolutely essential to production'. Fujio Cho of Toyota identified seven types of waste that must be eliminated from all operations processes. They are, waste from over-production, waste from waiting time, transportation waste, inventory waste, processing waste, waste of motion and waste from product defects. Beyond this, authorities on Toyota claim that its strength lies in understanding the differences between the tools and practices used with Toyota operations and the overall philosophy of their approach to lean synchronisation. This is what some have called the apparent paradox of the Toyota production system, 'namely, that activities, connections and production flows in a Toyota factory are rigidly scripted, yet at the same time Toyota's operations are enormously flexible and adaptable. Activities and processes are constantly being challenged and pushed to a higher level of performance, enabling the company to continually innovate and improve.' While come adopters of lean synchronisation principles may think they have 'done lean', Toyota simply changes the goal to constantly challenge improvement. As such, we see a key distinction between those that see lean as a specific end-point to be achieved through the application of a series of improvement tools, and those such as Toyota who treat lean as a philosophy that defines a way of conducting business

EXAMPLE

Kanban control at Torchbox web designers[2]

Torchbox is an independently owned web design and development company based in Oxfordshire (see the short case description in Chapter 1 'Torchbox: award-winning web designers' for more details). Despite relying largely on its staff's creative capabilities, this does not mean that lean techniques have no place in its operations. On the contrary, it makes full use of the 'kanban' approach to controlling its work as it progresses through the web design process. '*We know that kanban control originated from car manufacturers like Toyota, but our development teams can also benefit from its basic principles*' said Edward Kay, the Head of Production at Torchbox. '*It is a way of scheduling work based on what needs to be produced and what resources are available to produce it with. At Torchbox we use a large magnetic whiteboard (called the "kanban board") to track completed features through each stage of the design process; from discovery through development, design, demo, deployment and on to the finish of the design (called the "done" stage). Each feature has its own paper slip that physically moves across the board, held in place with a magnet. You can't have more features in progress than the number of magnets you have to hold them in place, so the principle is enforced with a physical limitation. When one feature enters the "Done" column, another one can be pulled through into discovery. There's a pulling process, where completing work where completing work allows you to start on something new.*'

At the start of every day, the team has a stand-up meeting at the kanban board where each member quickly runs through what they did the day before, and what they'll do in the coming day. Each developer has a few tokens that they place on the features they're working on. This helps link up the 'big picture' of how a design's features are developing, with the 'little picture' of what each developer is working on each day, and helps teams to make sure that all work being done is being tracked across the board.

'*One of the big benefits of using kanban*', said Edward Kay, '*is that because we're visualising the steps a feature goes through to be completed, we're able to see where the bottlenecks are that work gets held up on. If we're finding that a project's features keep getting held up in the design stage, we can bring more designers onto the project to widen the bottleneck. Using kanban with feature-driven development helps us constantly deliver value to our clients. This more measured and controlled approach to handling and controlling incoming work helps ensure that every hour we work produces an hour's worth of value. Ultimately, it's all about delivering great products on time and to budget, and kanban is a great tool to help achieve this.*'

What do the operations have in common?

Here are two types of operation separated by product, culture, size, location and their route to adopting lean synchronisation principles. Toyota took decades to develop a fully integrated and coherent philosophy to managing their operations and have become one of the world's leading and most profitable automotive companies as a result. Torchbox is far earlier in their path to lean synchronisation, yet they have adopted and adapted several ideas from the lean synchronisation philosophy and gained benefits. The exact interpretation of what 'lean' means in practice will differ, but it still has the potential to improve their service delivery. That is because, notwithstanding the differences between the two operations, the basic principles of lean synchronisation remain the same, namely aiming to achieve perfect synchronisation through smooth and even flow. But, lean synchronisation is an *aim*. It is not something that can simply be implemented overnight. Both these organisations have worked hard at overcoming the barriers to lean synchronisation. These can be summarised as, the elimination of all waste, the involvement of everyone in the business and the adoption of a continuous improvement philosophy. The focus on eliminating waste uses four important methods: streamlining flow; making sure that supply matches demand exactly; increasing process flexibility; and reducing the effects of variability.

Before further discussion it is important to be clear on the distinction between, what is the *aim* (lean synchronisation), the *approach to overcoming the barriers* to achieving lean synchronisation, the *methods of eliminating waste* and the various *techniques* that can be used to help eliminate waste. The relationship between these elements is shown in Figure 11.4.

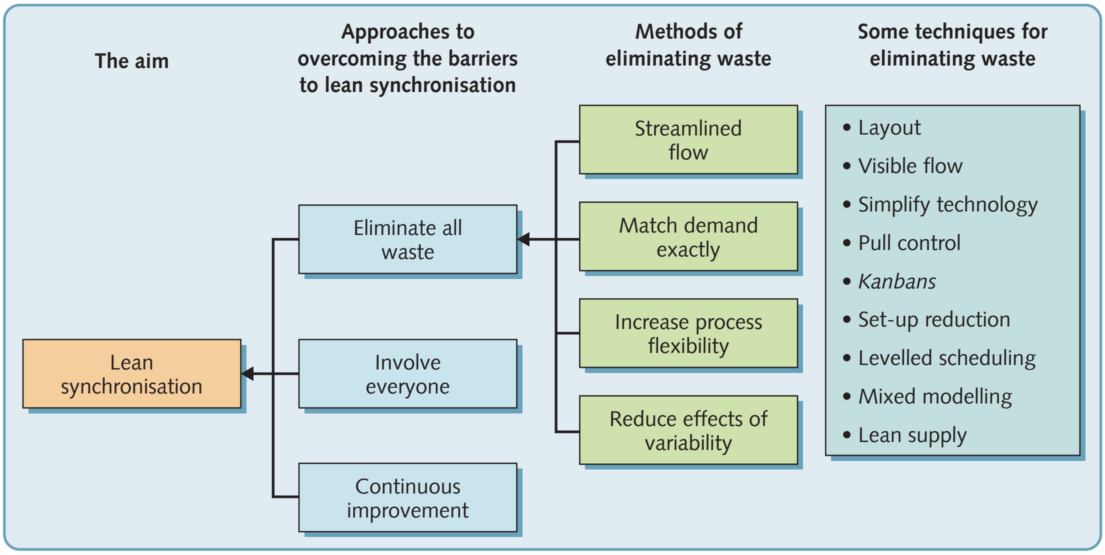

Figure 11.4 Schematic of the issues covered in this chapter

DIAGNOSTIC QUESTION

What are the barriers to 'lean synchronisation'?

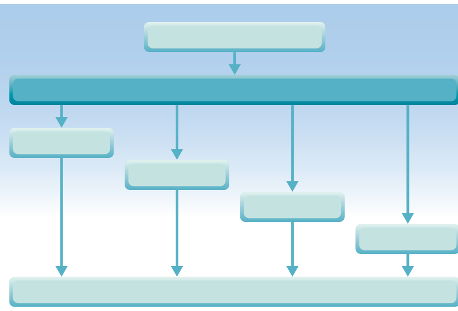

The aim of pure lean synchronisation represents an ideal of smooth, uninterrupted flow without delay, waste, or imperfection of any kind. The supply and demand between stages in each process, between processes in each operation and between operations in each supply network are all perfectly synchronised. It represents the ultimate in what customers are looking for from an operation. But first one must identify the barriers to achieving this ideal state. We group these under three headings:

1. A failure to eliminate waste in all parts of the operation
2. A failure to harness the contribution of all the people within the operation
3. A failure to establish improvement as a continuous activity

The waste elimination barrier

Arguably the most significant part of the lean philosophy is its focus on the elimination of all forms of waste. Waste can be defined as any activity that does not add value. For example, a study by Cummins, the engine company, showed that, at best, an engine was only being worked on for 15 per cent of the time it was in the factory.[3] At worst, this fell to 9 per cent, which meant that for 91 per cent of its time, the operation was adding cost to the engine, not adding value. Although already a relatively efficient manufacturer, the results alerted Cummins to the enormous waste which still lay dormant in its operations, and which no performance measure then in use had exposed. Cummins shifted its objectives to reducing the wasteful activities and to enriching the value-added ones. Exactly the same phenomenon applies in service processes. Relatively simple requests, such as applying for a driving licence, may only take a few minutes to actually process, yet take days (or weeks) to be returned.

OPERATIONS PRINCIPLE

Focusing on synchronous flow exposes sources of waste.

Identifying waste is the first step towards eliminating it. Toyota has described seven types. Here we consolidate these into four broad categories of waste that apply in many different types of operation

Waste from irregular flow

Perfect synchronisation means smooth and even flow through processes, operations and supply networks. Barriers that prevent streamlined flow include the following:

- *Waiting time:* machine efficiency and labour efficiency are two popular measures that are widely used to measure machine and labour waiting time, respectively. Less obvious is the time when products, customers or information wait as inventory or queues, there simply to keep operators busy.
- *Transport:* moving items or customers around the operation, together with double and triple handling, does not add value. Layout changes that bring processes closer together, improvements in transport methods and workplace organisation can all reduce waste.
- *Process inefficiencies:* the process itself may be a source of waste. Some operations may only exist because of poor component design, or poor maintenance, and so could be eliminated.
- *Inventory:* regardless of type (product, customer, information) all inventories should become a target for elimination. However, it is only by tackling the causes of inventory or queues, such as irregular flow, that it can be reduced.
- *Wasted motions:* an operator may look busy but sometimes no value is being added by the work. Simplification of work is a rich source of reduction in the waste of motion.

Waste from inexact supply

Perfect synchronisation also means supplying exactly what is wanted, exactly when it is needed. Any under- or over-supply and any early or late delivery will result in waste, something we have already explored in the capacity management chapter in particular. Barriers to achieving and exact match between supply and demand include the following.

- *Over-production or under-production:* supplying more than, or less than, is immediately needed by the next stage, process or operation. (This is the greatest source of waste according to Toyota.)
- *Early or late delivery:* items should only arrive exactly when they are needed. Early delivery is as wasteful as late delivery.
- *Inventory:* again, all inventories should become a target for elimination. However, it is only by tackling the causes of inventory, such as inexact supply, that it can be reduced.

Waste from inflexible response

Customer needs can vary, in terms of what they want, how much they want and when they want it. However, processes usually find it more convenient to change what they do relatively infrequently, because every change implies some kind of cost. That is why hospitals schedule specialist clinics only at particular times, and why machines often make a batch of similar products together. Yet responding to customer demands exactly and instantaneously requires a high degree of process flexibility. Symptoms of inadequate flexibility include the following:

- *Large batches:* sending batch of items through a process inevitably increases inventory as the batch moves through the whole process.
- *Delays between activities:* the longer the time (and the cost) of changing over from one activity to another, the more difficult it is to synchronise flow to match customer demand instantaneously.
- *More variation in activity mix than in customer demand:* if the mix of activities in different time periods varies more than customer demand varies, then some 'batching' of activities must be taking place.

Waste from variability

Synchronisation implies exact levels of quality. If there is variability in quality levels then customers will not consider themselves as being adequately supplied. Variability therefore is an important barrier to achieving synchronised supply. Symptoms of poor variability include the following:

- *Poor reliability of equipment:* unreliable equipment usually indicates a lack of conformance in quality levels. It also means that there will be irregularity in supplying customers. Either way, it prevents synchronisation of supply.
- *Defective products or services:* waste caused by poor quality is significant in most operations. Service or product errors cause both customers and processes to waste time until they are corrected.

Muda, Mura, Muri

An alternative, but common, classification of waste are the Japanese terms Muda, Mura and Muri:

- Muda – are activities in a process that are wasteful because they do not add value to the operation or the customer. The main causes of these wasteful activities are likely to be poorly communicated objectives (including not understanding the customer's requirements), or the inefficient use of resources. The implication of this is that for an activity is to be effective, it must be properly recorded and communicated to whoever is performing it.
- Mura – means 'lack of consistency' or unevenness that results in periodic overloading of staff or equipment. For example, if activities are not properly documented so that different people at different times perform a task differently, then not surprisingly, the result of the activity may be different. The negative effects of this are similar to a lack of dependability.
- Muri – means absurd or unreasonable. It is based on the idea that unnecessary or unreasonable requirements put on a process will result in poor outcomes. The implication of this is that appropriate skills, effective planning, accurate estimation of times and schedules will avoid this 'muri' overloading waste. In other words, waste can be caused by failing to carry out basic operations planning tasks such as prioritising activities (sequencing), understanding the necessary time (scheduling) and resources (loading) to perform activities. All these issues were discussed in Chapter 10.

These three causes of waste are obviously related. When a process is inconsistent (mura), it can lead to the overburdening of equipment and people (muri) that, in turn, will cause all kinds of non-value-adding activities (muda).

EXAMPLE

No more 'faffing around' in the kitchen[4]

For most people who work long hours, cooking is something they do not have time for. Knowing this, the celebrity chef Jamie Oliver has written a book *Jamie's 30-minute meals,* whose philosophy is that cooking a delicious dinner should be as quick as and cheaper than buying and heating a takeaway. The book presents 50 ready-made menus with 3 to 4 courses per menu, designed to take no more than 30 minutes to prepare. To achieve this performance Jamie has, perhaps inadvertently, applied the principles and methods of lean synchronisation to the everyday activity of cooking.

Let's imagine that your family is coming over for dinner and you want to surprise them with a new Indian multi-course meal with chicken, rice, salad on the side and of course a dessert. Traditionally, you would search and look up four different recipes, one for each dish. Because all recipes come from different places, you need to work out the quantity of food to buy, doing the maths in case of shared ingredients across the dishes, how to allocate pots, pans and other equipment to the different ingredients, and most importantly, you need to work out in what order to prepare things, especially if you want all your dishes ready at the same time. Jamie's approach significantly reduces this complexity by ensuring dishes are prepared right when the next step in the process needs it, regardless of which dish it is. In other words, dishes are not cooked in

sequence, one after another, but they are prepared and completed simultaneously.

If we identify all the tasks related to preparing the salad (e.g., chopping the vegetables) with the letter A, cooking the rice (e.g., blending) with letter B, cooking the chicken with letter C, and finally making the dessert with the letter D, then in the traditional way of cooking our task scheduling would look something like AAAA BBBBBBB CCCCCCC DDDD. This results in batching, waiting time and causing dishes to be ready before the dinner is supposed to be served. Conversely, Jamie Oliver's 30-minute cooking involves scheduling tasks in a sequence like ABCDACBAD-CBABDC, where single tasks related to different dishes follow smoothly, as the chef chops a salad ingredient, then blends the rice, then chops some more salad ingredients while the chicken is being roasted in the oven and a part of the desert is being prepared. This way, all dishes are ready at the same time, just in time and nothing is prepared before it has to be, avoiding any form of waste. Such a levelled approach to scheduling is called heijunka (mixed modelling) in lean synchronisation.

In addition, Jamie's lean cooking builds on reduced setup times. At the beginning of each recipe, the equipment needed to prepare the menu is presented under the headline 'To Start'. Other necessary preparations, such as heating the oven, are also specified. Having all equipment ready from the start saves time in the process, and is, according to Jamie, a prerequisite for getting done in 30 minutes. The use of simple equipment that is suitable for many different purposes also makes the process quicker as changeovers are minimised. The rationale is to make the most out of the time available, eliminating the 'faffing around' in cooking (non-value-added activity in OM language!) and leaving only what is strictly 'good, fast cooking', without compromising on quality.

Short-term pain for a long-term gain

A paradox in the lean synchronisation concept is that adoption may mean some sacrifice of capacity utilisation. In organisations that place a high value on the utilisation of capacity, this can prove particularly difficult to accept. But it is necessary. Return to the process shown in Figure 11.2. When stoppages occur in the traditional system, the buffers allow each stage to continue working and thus achieve high capacity utilisation. The high utilisation does not necessarily make the system as a whole produce more parts or process more customers. Often the extra processing goes into the large buffer inventories or queues. In a synchronised lean process, any stoppage will affect the rest of the system, causing stoppages throughout the operation. This will necessarily lead to lower capacity utilisation, at least in the short term. However, there is no point in processing products, services, or information for the sake of it. Unless the output is useful and enables the operation as a whole to complete all required tasks, there is no point in continuing to work. In fact, processing just to keep utilisation high is not only pointless, it is counter-productive, because the extra inventory or queues that are created merely serve to make improvements less likely. Figure 11.5 illustrates the two approaches to capacity utilisation.

OPERATIONS PRINCIPLE

Focusing on lean synchronisation can initially reduce resource utilisation.

The involvement barrier

An organisational culture that supports lean synchronisation must place a very significant emphasis on involving everyone in the organisation. This approach to people management (sometimes called the 'respect-for-people' system, after a rough translation from the Japanese) is seen by

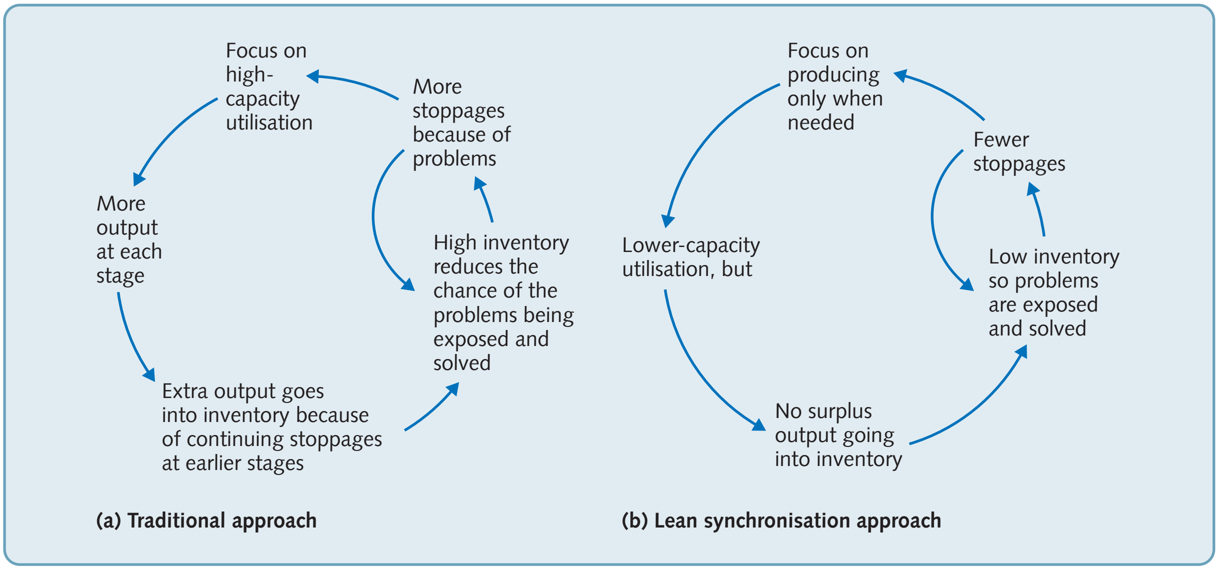

Figure 11.5 Lean The different views of capacity utilisation in (a) traditional and (b) lean synchronisation approaches to planning and controlling flow

some as the most controversial aspect of the lean philosophy. It encourages (and often requires) team-based problem solving, job enrichment, job rotation and multi-skilling. The intention is to encourage a high degree of personal responsibility, engagement and 'ownership' of the job. Some Japanese companies refer to operationalising of the 'involvement of everyone' principle by adopting 'basic working practices'. They are held to be the basic preparation of the operation and its employees for implementing lean synchronisation. They include the following:

- **Discipline** – Work standards that are critical for the safety of staff, the environment and quality must be followed by everyone all the time.
- **Flexibility** – It should be possible to expand responsibilities to the extent of people's capabilities. This applies as equally to managers as it does to shop-floor personnel. Barriers to flexibility, such as grading structures and restrictive practices, should be removed.
- **Equality** – Unfair and divisive personnel policies should be discarded. Many companies implement the egalitarian message through to company uniforms, consistent pay structures which do not differentiate between full-time staff and hourly-rated staff and open-plan offices.
- **Autonomy** – Delegate responsibility to people involved in direct activities so that management's task becomes one of supporting processes. Delegation includes giving staff the responsibility for stopping processes in the event of problems, scheduling work, gathering performance monitoring data and general problem solving.
- **Development of personnel** – Over time, the aim is to create more company members who can support the rigours of being competitive.
- **Quality of working life (QWL)** – This may include, for example, involvement in decision-making, security of employment, enjoyment and working area facilities.
- **Creativity** – This is one of the indispensable elements of motivation. Creativity in this context means not just doing a job, but also improving how it is done and building the improvement into the process.
- **Total people involvement** – Staff take on more responsibility to use their abilities to the benefit of the company as a whole. They are expected to participate in activities such as the selection of new recruits, dealing directly with suppliers and customers over schedules, quality issues, and delivery information, spending improvement budgets and planning and reviewing work done each day through communication meetings.

The concept of continuous learning is also central to the 'involvement of everyone' principle. For example, Toyota's approach to involving its employees includes using a learning method that allows employees to discover the Toyota production system rules through problem solving. So, while the job is being performed, a supervisor/trainer asks a series of questions that gives the employee deeper insights into the work.[5] These questions could be:

- How do you do this work?
- How do you know you are doing this work correctly?
- How do you know that the outcome is free of defects?
- What do you do if you have a problem?

The continuous improvement barrier

Lean synchronisation objectives are often expressed as ideals, such as our previous definition: 'to meet demand instantaneously with perfect quality and no waste'. While any operation's current performance may be far removed from such ideals, a fundamental lean belief is that it is possible to get closer to them over time. Without such beliefs to drive progress, lean proponents claim improvement is more likely to be transitory than continuous. This is why the concept of continuous improvement is such an important part of the lean philosophy. If its aims are set in terms of ideals which individual organisations may never fully achieve, then the emphasis must be on the way in which an organisation moves closer to the ideal state. The Japanese word that incorporates the idea of continuous improvement is 'kaizen'. It is one of the main pillars of process improvement and is explained fully in Chapter 12.

Techniques to address the four sources of waste

Of the three barriers to achieving lean synchronisation (reduce waste, involve everyone and adopt continuous improvement), the last two are addressed further in Chapter 12. Therefore the rest of this chapter is devoted to what could be called the 'core of lean synchronisation'. These are a collection of 'just-in-time' tools and techniques that are the means of cutting out waste. Although many of these techniques are used to reduce waste generally within processes, operations and supply networks, as earlier, we will group the approaches to reducing waste under the four main headings: streamlining flow; matching demand exactly; increasing process flexibility; and reducing the effects of variability.

DIAGNOSTIC QUESTION

Is flow streamlined?

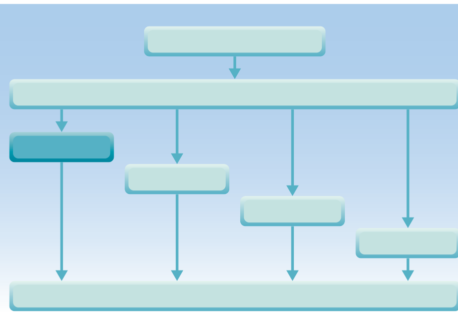

The smooth flow of materials, information and people in the operation is a central idea of lean synchronisation. Long process routes provide opportunities for delay and inventory / queue build-up, add no value and slow down throughput time. So, the first contribution any operation can make to streamline flow is to reconsider the basic layout of its processes. Primarily, reconfiguring the layout of a process to aid lean synchronisation involves moving it down the 'natural diagonal' of process design that was discussed in Chapter 5. Broadly speaking, this means moving from functional layouts towards cell-based layouts, or from cell-based layouts towards line (or product) layouts. Either way, it is necessary to move towards a layout that brings more systematisation and control to the process flow. At a more detailed level, typical layout techniques include: placing workstations close together so that inventory or queues physically just cannot build-up because there is no space for it to do so; and arranging workstations in such a way that all those who contribute to a common activity are in sight of each other and can provide mutual help. For example, at the Virginia Mason Medical Centre, Seattle, USA, a leading proponent of lean synchronisation in healthcare, many of the waiting rooms have been significantly reduced in their capacity or removed entirely. This forces a focus on the flow of the whole process because patients have literally nowhere to be stored.

OPERATIONS PRINCIPLE

Simple, transparent flow exposes sources of waste.

Examine the shape of process flow

The pattern that flow makes within or between processes is not a trivial issue. Processes that have adopted the practice of curving line arrangements into U-shaped or 'serpentine' arrangements can have a number of advantages (U-shapes are usually used for shorter lines and serpentines for longer lines). One authority[6] sees the advantages of this type of flow patterns as *staffing flexibility and balance,* because the U-shape enables one person to tend several activities; *rework,* because it is easy to return faulty work to an earlier station; *free flow,* because long straight lines interfere with cross travel in the rest of the operation; and *teamwork,* because the shape encourages a team feeling.

Ensure visibility

Appropriate layout also includes the extent to which all movement is transparent to everyone within the process. High visibility of flow makes it easier to recognise potential improvements to flow. It also promotes quality within the process because the more transparent the operation or process, the easier it is for all staff to share in its management and improvement. Problems are more easily detectable and information becomes simple, fast and visual. Visibility measures include the following:

- clearly indicated process routes using signage
- performance measures clearly displayed in the workplace
- coloured lights used to indicate stoppages
- an area is devoted to displaying samples of one's own and competitors' process outputs, together with samples of good and defective output
- visual control systems (e.g. kanbans, discussed later).

An important technique used to ensure flow visibility is the use of simple, but highly visual signals to indicate that a problem has occurred, together with operational authority to stop the process. For example, on an assembly line, if an employee detects some kind of quality problem, he or she could activate a signal that illuminates a light (called an 'andon' light) above the workstation and stops the line. Although this may seem to reduce the efficiency of the line, the idea is that this loss of efficiency in the short term is less than the accumulated losses of allowing defects to continue on in the process. Unless problems are tackled immediately, they may never be corrected.

EXAMPLE

Andons in Amazon[7]

The principles of an andon cord or light have been applied in Amazon. Every day, service agents at Amazon receive calls from customers who are unhappy with some aspect of the product delivered to them. Customer agents dealing with these complaints are now empowered to make judgements on the extent to which such complaints may be systemic. In cases

where they suspect it's a repetitive defect, service agents can 'stop the line' for a particular product. This involves taking the product off the website while the problem is fully investigated. According to Amazon, the improved visibility of the system has eliminated tens of thousands of defects a year and has also given service agents a strongly sense of being able to deal effectively with customer complaints. Now an agent can not only refund the individual customer, they can also tell the customer that others won't receive products until the problem has been properly investigated. The firm also claims that around 98 per cent of the times when the andon cord is pulled, there really is a systemic problem, highlighting the value of trusting in their service agents to make sensible decisions on when to and when not to stop the line.

Use small-scale simple process technology

There may also be possibilities to encourage smooth streamlined flow through the use of small-scale technologies; that is, using several small units of process technology, rather than one large unit. In a component manufacturer, for example, small machines have several advantages over large ones. First, they can process different products and services simultaneously. For example, in Figure 11.6 one large machine produces a batch of A, followed by a batch of B, and followed by a batch of C. However, if three smaller machines are used, they can each produce A, B, or C simultaneously. The system is also more robust. If one large machine breaks down, the whole system ceases to operate. If one of the three smaller machines breaks down, it is still operating at two-thirds effectiveness. Small machines are also easily moved, so that layout flexibility is enhanced, and the risks of making errors in investment decisions are reduced. However, investment in capacity may increase in total because parallel facilities are needed, so utilisation may be lower (see the earlier arguments).

Examining all elements of throughput time

Throughput time is often taken as a surrogate measure for waste in a process. The longer that items being processed are held in inventory, moved, checked, or subject to anything else that does not add value, the longer they take to progress through the process. So, looking at exactly what happens to items within a process is an excellent method of identifying sources of waste.

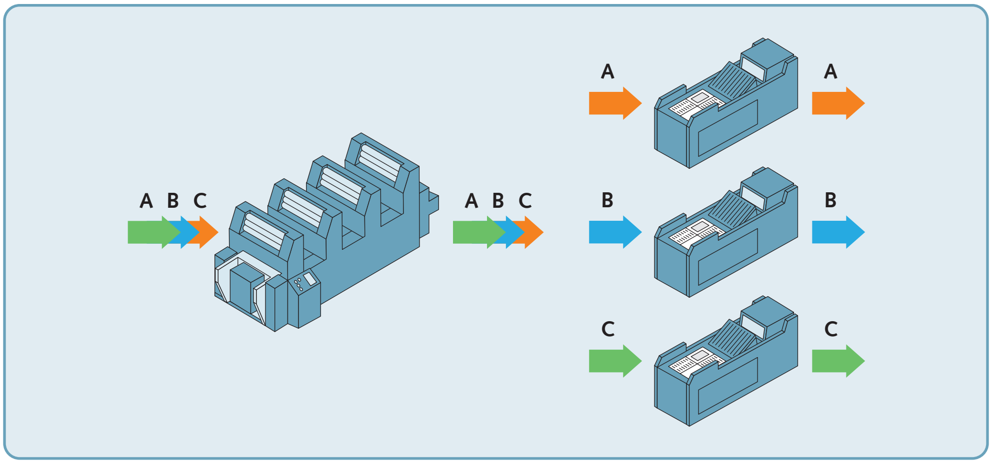

Figure 11.6 **Using several small machines rather than one large one, allows simultaneous processing, is more robust and is more flexible.**

Value stream mapping (also known as 'end-to-end' system mapping) is a simple but effective approach to understanding the flow of material and information as a product or service has value added as it progresses through a process, operation, or supply chain. It visually maps a product or service's 'production' path from start to finish. In doing so, it records, not only the direct activities of creating products and services, but also the 'indirect' information systems that support the direct process. It is called 'value stream' mapping because it focuses on value-adding activities and distinguishes between value-adding and non-value-adding activities. It is similar to process mapping (see Chapter 6) but different in four ways:

- It uses a broader range of information than most process maps.
- It is usually at a higher level (5–10 activities) than most process maps.
- It often has a wider scope, frequently spanning the whole supply chain.
- It can be used to identify where to focus future improvement activities.

A value stream perspective involves working on (and improving) the 'big picture', rather than just optimising individual processes. Value stream mapping is seen by many practitioners as a starting point to help recognise waste and identify its causes. It is a four-step technique that identifies waste and suggests ways in which activities can be streamlined. First, it involves identifying the value stream (the process, operation or supply chain) to map. Second, it involves physically mapping a process, then above it mapping the information flow that enables the process to occur. This is the, so-called 'current state' map. Third, problems are diagnosed and changes suggested, making a future state map that represents the improved process, operation or supply chain. Finally, the changes are implemented. Figure 11.7 shows a value stream map for an industrial air conditioning installation service. The service process itself is broken down into five relatively large stages and various items of data for each stage are marked on the chart. The type of data collected here does vary, but all types of value stream maps compare the total throughput time with the amount of value-added time within the larger process. In this case, only 8 of the 258 hours of the process are value-adding.

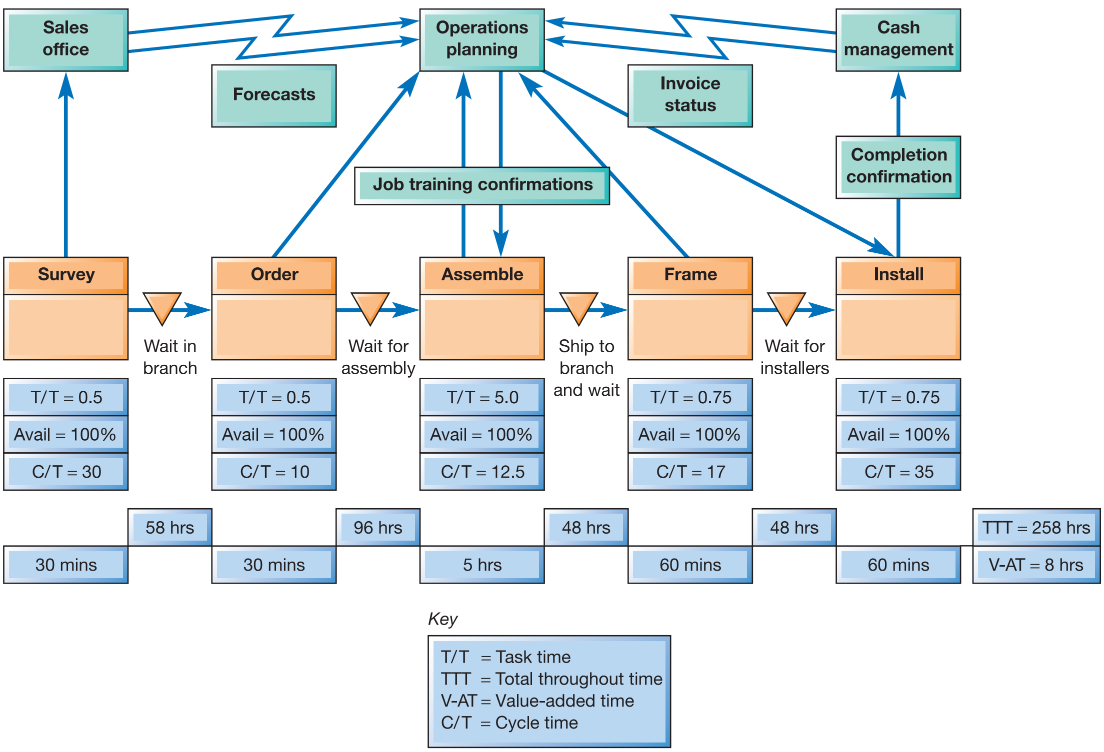

Figure 11.7 **Value stream map for an industrial air conditioning installation service**

DIAGNOSTIC QUESTION

Does supply exactly match demand?

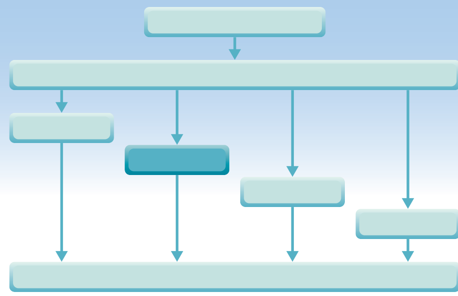

The value of the supply of products or services is always time dependent. Something that is delivered early or late often has less value than something that is delivered exactly when it is needed. We can see many everyday examples of this. For example, parcel delivery companies charge more for guaranteed faster delivery; Airlines charge extra for priority boarding; visa applications centres typically offer rapid turn-around services; and private hospitals charge more by offering faster access to their services. The closer to instantaneous delivery we can get, the more value the delivery has for us and the more we are willing to pay for it. In fact, delivery earlier than it is required can in some case be just as harmful as late delivery because it results in information inventories that serve to confuse flow through the process. For example, an Australian tax office used to receive applications by mail, open the mail and send it through to the relevant department who, after processing it, sent it to the next department. This led to piles of unprocessed applications building up within its processes causing problems in tracing applications, and losing them, sorting through and prioritising applications, and worst of all, long throughput times. Now they only open mail when the stages in front can process it. Each department requests more work only when they have processed previous work. Likewise, Tesco, one of Europe's leading food retailers has financial penalties for both late *and* early deliveries into its cross-docking distribution centres!

OPERATIONS PRINCIPLE

Delivering only and exactly, what is needed and when it is needed, smooths flow and exposes waste.

Pull control

The exact matching of supply and demand is often best served by using 'pull control' wherever possible (discussed in Chapter 10). At its simplest, consider how some fast-food restaurants cook and assemble food and place it in the warm area only when the customer-facing server has sold an item. Production is being triggered only by real customer demand. Similarly, supermarkets usually replenish their shelves only when customers have taken sufficient products off the shelf. The movement of goods from the 'back office' store to the shelf is triggered only by the 'empty-shelf' demand signal. Some construction companies make it a rule to call for material deliveries to its sites, only the day before those items are actually needed. This not only reduces clutter and the chances of theft, it speeds up throughput time and reduces confusion and inventories. The essence of pull control is to let the downstream stage in a process, operation, or supply network, pull items through the system rather than have them 'pushed' to them by the supplying stage.

Kanbans

The use of kanbans is one method of operationalising pull control. Kanban is the Japanese for card or signal. It is sometimes called the 'invisible conveyor' that controls the transfer of items between the stages of an operation. In its simplest form, it is a card used by a customer stage to instruct its supplier stage to send more items. Kanbans can also take other forms. In some Japanese companies, they are solid plastic markers or even coloured Ping-Pong balls. Whichever kind of kanban is being used, the principle is always the same; the receipt of a kanban triggers the movement, production or supply of one unit or a standard container of units. If two kanbans are received, this triggers the movement, production or supply of two units or standard containers of units, and so on. Kanbans are the only means by which movement, production or supply can be authorised. Some companies use 'kanban squares'. These are marked spaces on the shop floor or bench that are drawn to fit one or more work pieces or containers. Only the existence of an empty square triggers production at the stage that supplies the square. Likewise 'kanban' whiteboard are increasingly used to 'pull' activity through the service process (see the Torchbox example).

DIAGNOSTIC QUESTION

Are processes flexible?

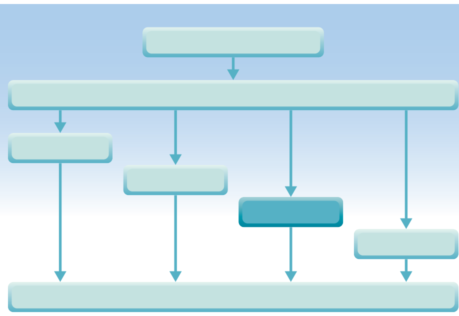

Responding exactly and instantaneously to customer demand implies that operations resources need to be sufficiently flexible to change both what they do and how much they do of it, without incurring high cost or long delays. In fact, flexible processes (often with flexible technologies) can significantly enhance smooth and synchronised flow. For example, new publishing technologies allow professors to assemble printed and e-learning course material customised to the needs of individual courses or even individual students. In this case flexibility is allowing customised, small batches to be delivered 'to order'. In another example, a firm of lawyers used to take ten days to prepare its bills for customers. This meant that customers were not asked to pay until ten days after the work had been done. Now they use a system that everyday updates each customer's account. So, when a bill is sent it includes all work up to the day before the billing date. The principle here is that process inflexibility also delays cash flow.

OPERATIONS PRINCIPLE

Changeover flexibility reduces waste and smooths flow.

Reduce setup times

For many technologies, increasing process flexibility, means reducing setup times; defined as the time taken to change over the process from one activity to the next. Compare the time it takes you to change the tyre on your car with the time taken by a Formula 1 team. Setup reduction can be achieved by a variety of methods such as cutting out time taken to search for tools and equipment, the pre-preparation of tasks which delay changeovers, and the constant practice of setup routines. The other common approach to setup time reduction is to convert work that was previously performed while the process was stopped (called *internal* work) to work that is performed while the process is running (called *external* work).

In a manufacturing context, making this transition from *internal* to *external* wok may include making equipment capable of performing all required tasks so that changeovers become a simple adjustment, and using simple devices to facilitate change of equipment – for example, having equipment on casters to reconfigure processes quickly. In a service context, reducing setup times is also critical. For example, airlines can't make money from aircraft that are sitting idle on the ground – called 'running the aircraft hot' in the industry. For many smaller airlines, the biggest barrier to running hot is that their markets are not large enough to justify passenger flights during the day *and* night. So, in order to avoid aircraft being idle over night, they must be used in some other way. That was the motive behind Boeing's 737 'Quick Change' (QC) aircraft. With it, airlines have the flexibility to use it for passenger flights during the day and, with less than a one hour changeover (setup) time, use it as a cargo airplane throughout the night. Boeing engineers designed frames that hold entire rows of seats that could smoothly glide on and off the aircraft, allowing 12 seats to be rolled into place at once. When used for cargo, the seats are simply rolled out and replaced by special cargo containers designed to fit the curve of the fuselage and prevent damage to the interior. Before reinstalling the seats the sidewalls are thoroughly cleaned so that once the seats are in place, passengers cannot tell the difference between a QC aircraft and a normal 737. Airlines like Aloha Airlines, which serves Hawaii, particularly value the aircraft's flexibility. It allows them to provide frequent reliable services in both passenger and cargo markets. So the aircraft that has been carrying passengers around the islands during the day can be used to ship fresh supplies over night to the hotels that underpin the tourist industry.

DIAGNOSTIC QUESTION

Is variability minimised?

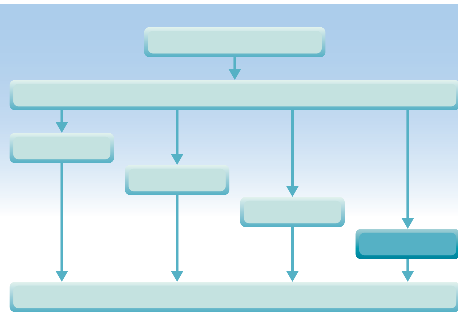

One of the biggest causes of the variability that will disrupt flow and prevent lean synchronisation is variation in the quality of items. This is why a discussion of lean synchronisation should always include an evaluation of how quality conformance is ensured within processes. In particular, the principles of statistical process control (SPC) can be used to understand quality variability. Chapter 13 and its supplement on SPC examine this subject, so in this section we shall focus on other causes of variability. The first of these is variability in the mix of products and services moving through processes, operations, or supply networks.

OPERATIONS PRINCIPLE

Variability, in product/service quality, or quantity, or timing, acts against smooth flow and waste elimination.

Level schedules as much as possible

Levelled scheduling (or *heijunka*) means keeping the mix and volume of flow between stages even over time. For example, instead of producing 500 parts in one batch, which would cover the needs for the next three months, levelled scheduling would require the process to make only one piece per hour regularly. Thus, the principle of levelled scheduling is very straightforward; however, the requirements to put it into practice are quite severe, although the benefits resulting from it can be substantial. The move from conventional to levelled scheduling is illustrated in Figure 11.7. Conventionally, if a mix of products or service were required in a time period (usually a month), a batch size would be calculated for each product and the batches produced in some sequence. Figure 11.8(a) shows three products that are produced in a 20-day time period in a production unit.

Quantity of product/service A required = 3,000
Quantity of product/service B required = 1,000
Quantity of product/service C required = 1,000
Batch size of product/service A = 600
Batch size of product/service B = 200
Batch size of product/service C = 200

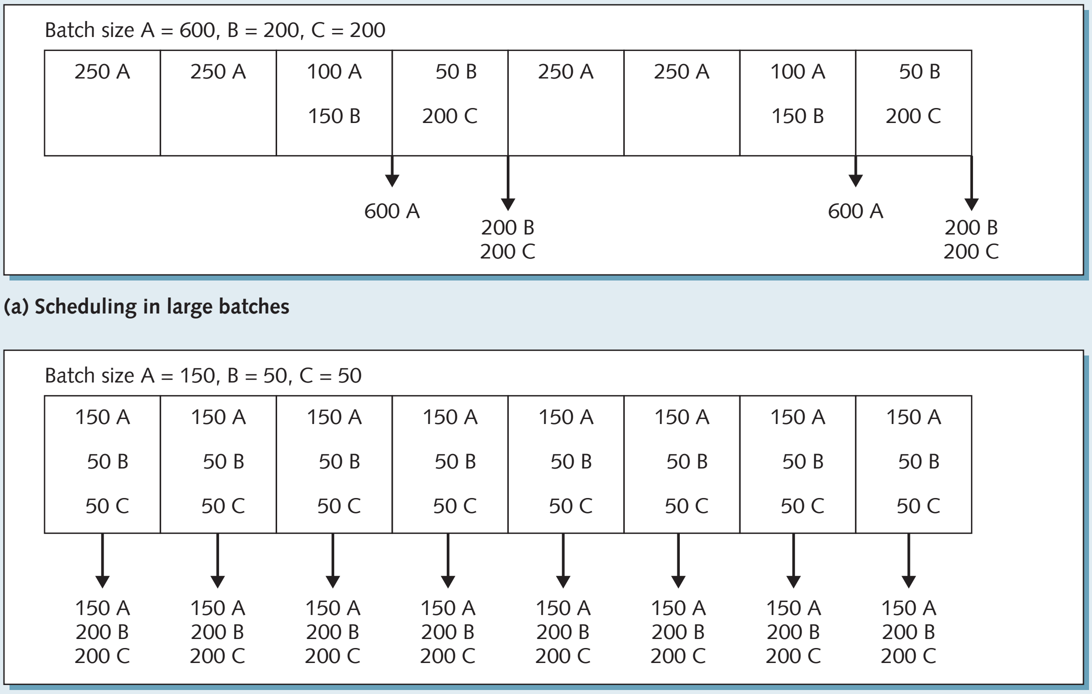

(a) Scheduling in large batches

(b) Levelled scheduling

Figure 11.8 **Levelled scheduling equalises the mix of products/service delivered each day**

Starting at day 1, the unit commences delivering product/service A. During day 3, the batch of 600 As is finished and dispatched to the next stage. The batch of Bs is started but is not finished until day 4. The remainder of day 4 is spent making the batch of Cs and both batches are dispatched at the end of that day. The cycle then repeats itself. The consequence of using large batches is, first, that relatively large amounts of inventory (or queues if this is a service setting and we're batching, for example different patient groups in an out-patient clinic) accumulate within and between the units, and second, that most days are different from one another in terms of what they are expected to produce (in more complex circumstances, no two days would be the same).

Now suppose that the flexibility of the unit could be increased to the point where the batch sizes for the products or service were reduced to a quarter of their previous levels without loss of capacity (see Figure 11.8(b):

Batch size of product/service A = 150
Batch size of product/service B = 50
Batch size of product/service C = 50

A batch of each product or service can now be completed in a single day, at the end of which the three batches are dispatched to their next stage. This will reduce the overall level of work-in-progress in the operation. Just as significant, however, is the effect on the regularity and rhythm of the work in each unit. Now every day in the month is the same in terms of what needs to be produced. This makes planning and control of each stage in the operation much easier. For example, if on day 1 of the month the daily batch of As was finished by 11.00 a.m., and all the batches were successfully completed in the day, then the following day the unit will know that, if it again completes all the As by 11.00 a.m., it is on schedule. When every day is different, the simple question 'Are we on schedule to complete our production today?' requires some investigation before it can be answered. However, when every day is the same, everyone in the unit can tell whether production is on target by looking at the clock. Control becomes visible and transparent to all, and the advantages of regular, daily schedules can be passed to upstream suppliers.

Level delivery schedules

A similar concept to levelled scheduling can be applied to many transportation processes. For example, a chain of convenience stores may need to make deliveries of all the different types of products it sells every week. Traditionally, it may have dispatched a truck loaded with one particular product around all its stores so that each store received the appropriate amount of the product that would last them for one week. This is equivalent to the large batches discussed in the previous example. An alternative would be to dispatch smaller quantities of all products in a single truck more frequently. Then, each store would receive smaller deliveries more frequently, inventory levels would be lower and the system could respond to trends in demand more readily because more deliveries means more opportunity to change the quantity delivered to a store. This is illustrated in Figure 11.9.

Adopt mixed modelling where possible

The principle of levelled scheduling can be taken further to give mixed modelling; that is, a repeated mix of outputs. Suppose that the machines in the production unit or employees in a service operation can be made so flexible that they achieve the JIT ideal of a batch size of one. The sequence of individual products or service emerging from the unit could be reduced progressively as illustrated in Figure 11.10. This would create a steady stream flowing continuously from the unit. However, the sequence does not always fall as conveniently as in Figure 11.10. The working times needed for different products or services are rarely identical and the ratios of required volumes are less convenient. For example, a small business tax return process is required to deal with different tax returns A (sole traders), B (partnerships) and C (limited companies with zero 0–3 employees) in the ratio 8:5:4. It could process 800 of A, followed by 500 of B, followed by 400 of A; or 80A, 50B and 40C. But ideally, to sequence the work as smoothly as (theoretically) possible, it would process in the order BACABACABACABACAB . . . repeated . . . repeated . . . etc. Doing this achieves relatively smooth flow (but does rely on significant process flexibility).

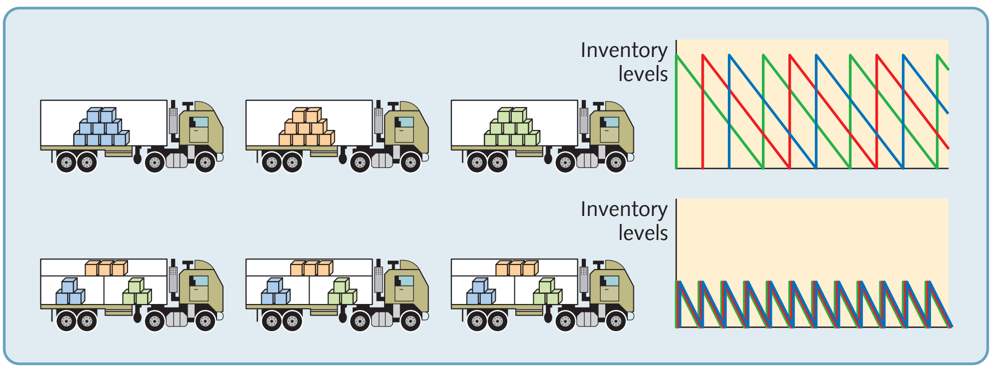

Figure 11.9 Delivering smaller quantities more often can reduce inventory or queue levels

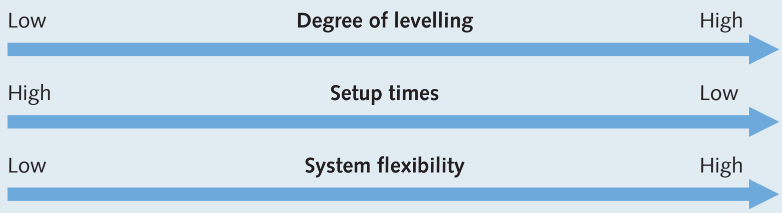

Large batches, e.g.	Small batches, e.g.	Mixed modelling, e.g.
200 A 120 B 80 C	5 A 3 B 2 C	A A B A B C A B C A

Figure 11.10 Levelled scheduling and mixed modelling: mixed modelling becomes possible as the batch size approaches one

Adopt total productive maintenance

Total productive maintenance (TPM) aims to eliminate the variability in operations processes caused by the effect of breakdowns. This is achieved by involving everyone in the search for maintenance improvements. Process owners are encouraged to assume ownership of their equipment and to undertake routine maintenance and simple repair tasks. These principles apply equally to service operations. For example, at a car wash, service employees regularly maintain their power-hoses to prevent unnecessary downtime, while university employees may be encouraged to regularly 'clean' e-mail inboxes, delete old files on their computers and update software with the aim of maintaining system availability speed. By doing so, maintenance specialists can then be freed to develop higher-order skills for improved maintenance systems. TPM is treated in more detail in Chapter 14.

DIAGNOSTIC QUESTION

Is lean synchronisation applied throughout the supply network?

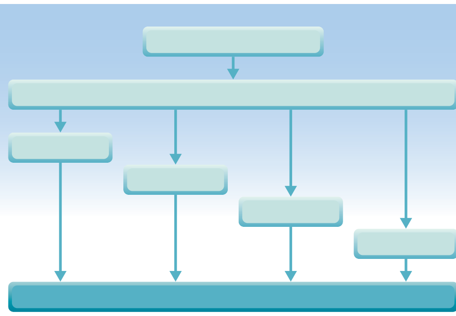

Although most of the concepts and techniques discussed in this chapter are devoted to the management of stages *within* processes and processes *within* an operation, the same principles can apply to the whole supply chain. In this context, the stages in a process are the whole businesses, operations or processes between which products flow. And as any business starts to approach lean synchronisation it will eventually come up against the constraints imposed by the lack of lean synchronisation of the other operations in its supply chain. So, achieving further gains must involve trying to spread lean synchronisation practice outward to its partners in the chain. Ensuring lean synchronisation throughout an entire supply network is clearly a far more demanding task than doing the same within a single process. It is a complex task. And it becomes more complex as more of the supply chain embraces the lean philosophy. The nature of the interaction between whole operations is far more complex than between individual stages within a process. A far more complex mix of products and services is likely to be being provided

OPERATIONS PRINCIPLE

The advantages of lean synchronisation apply at the level of the process, the operation and the supply network.

and the whole network is likely to be subject to a less predictable set of potentially disrupt events. To make a supply chain adopt lean synchronisation means more than making each operation in the chain lean. A collection of localised lean operations rarely leads to an overall lean chain. Rather one needs to apply the lean synchronisation philosophy to the supply chain as a whole. Yet the advantages from truly lean chains can be significant. Essentially, the principles of lean synchronisation are the same for a supply chain as they are for a process. Fast throughput throughout the whole supply network is still valuable and will save costs throughout the supply network. Lower levels of inventory (product, customer, or information) will still make it easier to achieve lean synchronisation. Waste is just as evident (and even larger) at the level of the supply network and reducing waste is therefore a critical task. Streamline flow, exact matching of supply and demand, enhanced flexibility and minimisng variability are all still tasks that will benefit the whole network. The principles of pull control can work between whole operations in the same way as they can between stages within a single process. In fact, the principles and the techniques of lean synchronisation are essentially the same no matter what level of analysis is being used. And because lean synchronisation is being implemented on a larger scale, and though the barriers are likely to be greater, so too will the potential benefits.

One of the weaknesses of lean synchronisation principles is that they are difficult to achieve when conditions are subject to unexpected disturbance (see the critical commentary later). This is especially a problem when applying lean synchronisation principles in the context of the whole supply network. Whereas unexpected fluctuations and disturbances do occur within operations, local management has a reasonable degree of control that it can exert in order to reduce them. Outside the operation, within the supply network, fluctuations can also be controlled to some extent (see Chapter 7), but it is far more difficult to do so. Nevertheless, it is generally held that, although the task is more difficult and although it may take longer to achieve, the aim of lean synchronisation is just as valuable for the supply network as a whole as it is for an individual operation.

Lean supply chains are like an air traffic control system[8]

The concept of the lean supply chain has been likened to an air traffic control system, in that it attempts to provide continuous, 'real-time visibility and control' to all elements in the chain. This is the secret of how the world's busiest airports handle thousands of departures and arrivals daily. All aircraft are given an identification number that shows up on a radar map. Aircraft approaching an airport are detected by the radar and contacted using radio. The control tower precisely positions the aircraft in an approach pattern, which it coordinates. The radar detects any small adjustments that are necessary, which are communicated to the aircraft. This real-time visibility and control can optimise airport throughput while maintaining extremely high safety and reliability.

Contrast this to how most supply chains are coordinated. Information is captured only periodically, probably once a day (or often even less frequently), and any adjustments to logistics, output levels at the various operations in the supply chain are adjusted, and plans rearranged. But imagine what would happen if this was how the airport operated, with only a 'radar snapshot' once a day. Coordinating aircraft with sufficient tolerance to arrange take-offs and landings every two minutes would be out of the question. Aircraft would be jeopardised or, alternatively, if aircraft were spaced further apart to maintain safety, throughput would be drastically reduced. Yet this is how most supply chains have traditionally operated. They use a daily 'snapshot' from their ERP systems (see Chapter 10 for an explanation of ERP). This limited visibility means operations must either space their work out to avoid 'collisions' (i.e. missed customer orders) thereby reducing output, or they must 'fly blind' thereby jeopardising reliability.

Lean and agile

One continuing debate on how lean principles can be applied across the supply chain concerns whether the supply network should be lean or 'agile'. Professor Martin Christopher of Cranfield University defines 'agility' as 'rapid strategic and operational adaptation to large-scale, unpredictable changes in the business environment. Agility implies responsiveness from one end of the supply chain to the other. It focuses upon eliminating the barriers to quick response, be they organisational or technical.' Other definitions stress that agility is the capability of operating profitably in a competitive environment of continually changing customer opportunities. The clue lies in how the word 'agile' is often defined; it implies being responsive, quick moving, flexible, nimble, active and constantly ready to change. But some proponents of operational agility go further than this. They see agility as also implying a rejection of a planning paradigm that makes any assumption of a predictable future. Like lean, it is more of a philosophy than an approach. Agile encourages a better match to what customers want by placing an emphasis on producing 'emergent' demand as opposed to rigid plans or schedules. Furthermore, rather than uncertainty and change being seen as things to be 'coped with' or preferably avoided, it should be embraced so that agility becomes changing faster than one's customer. Even less ambitious approaches to agility see it as more than simply organisational flexibility. It involves an organisational mastery of uncertainty and change, where people within the organisation, their capacity to learn from change and their collective knowledge are regarded as the organisation's greatest assets because they allow the operation to respond effectively to uncertainty and change. Continually inventing innovative business processes solutions to new market demands becomes a key operations objective.

All this seems very different to the underlying assumptions of the lean philosophy. Again, look at the word, lean means thin, having no superfluous fat. Lean attempts to eliminate waste and provide value to the customer throughout the entire supply chain. It thrives on standardisation, stability, defined processes and repeatability – not at all the way agility has been described. Lean is also a well-defined (although frequently misunderstood) concept. Agility, on the other hand, is a far newer and less 'operationalised' set of relatively strategic objectives. But some operational level distinctions can be inferred.

The type of principles needed to support a lean philosophy include such things as, simple processes, waste elimination, simple (if any) IT, the use of manual and robust planning and control as well as pull control and kanbans with overall MRP. Agile philosophies, by contrast, require effective demand management to keep close to market needs, a focus on customer relationship management, responsive supply coordination and visibility across the extended supply chain, continuous rescheduling and quick response in response to changing demand, short planning cycles, integrated knowledge management and fully exploited e-commerce solutions.

So are lean and agile philosophies fundamentally opposed? Well, yes and no. Certainly they have differing emphases. Saying that lean equals synchronised, regular flow and low inventory, and agile equals responsiveness, flexibility and fast delivery, may be something of a simplification, but it more or less captures the distinction between the two. But because they have different objectives and approaches does not mean that they cannot co-exist. Nor does it mean that there is a 'lean versus agile' argument to be resolved. The two approaches may not be complimentary, as some consultants claim, but both do belong to the general collection of methodologies that are available to help companies meet the requirements of their markets. In the same way as it was wrong to think that JIT would replace MRP, so 'agile' is not a substitute for lean.[9]

However, agile and lean are each more appropriate for differing market and product/service conditions. Put simply, if product/service variety or complexity is high, and demand predictability low then you have the conditions in which agile principles keep an operation ready to cope with instability in the business environment. Conversely, if product/service variety is low, and demand predictability high, then a lean approach can exploit the stable environment to achieve cost efficiency and dependability. So the two factors of product/service variety or complexity and demand uncertainty influence whether agile or lean principles should dominate. But what of the conditions where complexity and uncertainty are not related in this manner?

- When complexity is low, and demand uncertainty is also low (operations that produce commodities) lean planning and control is appropriate.
- When is complexity is low, and demand uncertainty is high (operations that produce fashion-based products/services) agile planning and control is appropriate.
- When complexity is high, and demand uncertainty is also high (operations that produce 'super value' products/services) project or requirements planning and control (for example of MRPII, see Chapter 10) is appropriate.
- When product/service complexity is high, and demand uncertainty is low (operations that produce 'consumer durable type' products/services) a combination of agile and lean planning and control is appropriate.

This last category has been rather clumsily called 'leagile'. Leagile is based on the idea that both lean and agile practices can be employed within supply chains. It envisages an inventory decoupling point that is the separation between the responsive (and therefore agile) 'front end' of the supply chain that reacts fast and flexibly to customer demand, and the efficient (and therefore lean) 'back-end' of the supply chain. This is not a new idea and in product-based supply chains involves 'making to *forecast*' before the decoupling point and 'making (or assembling, adapting or finishing) to *order'* after it. The idea has many similarities with the idea of 'mass customisation'. In a service context, work that is not for a specific customer, for example a consultancy company creating a 'platform' sector brief, may be worked on in a lean manner, because demand is relatively stable and predictable. By contrast, higher levels of customisation may then be required to adapt such a report to the specific needs of a client, and this may be carried out in an agile manner in the face of changing and often more unpredictable requests from the client.

Critical commentary

Lean synchronisation principles can be taken to an extreme. When just-in-time ideas first started to have an impact on operations practice in the West, some authorities advocated the reduction of between-process inventories of queues of customers to zero. While in the long term this provides the ultimate in motivation for operations managers to ensure the efficiency and reliability of each process stage, it does not admit the possibility of some processes always being intrinsically less than totally reliable. An alternative view is to allow inventories (albeit small ones) around process stages with higher than average uncertainty. This at least allows some protection for the rest of the system. The same ideas apply to just-in-time delivery between factories. The Toyota Motor Corp., often seen as the epitome of modern JIT, has suffered from its low inter-plant inventory policies. Both the Kobe earthquake and fires in supplier plants have caused production at Toyota's main factories to close down for several days because of a shortage of key parts. Even in the best-regulated manufacturing networks, one cannot always account for such events.

- One of the most counterintuitive issues in lean synchronisation is the way it appears to downplay the idea of capacity underutilisation. And it is true that, when moving

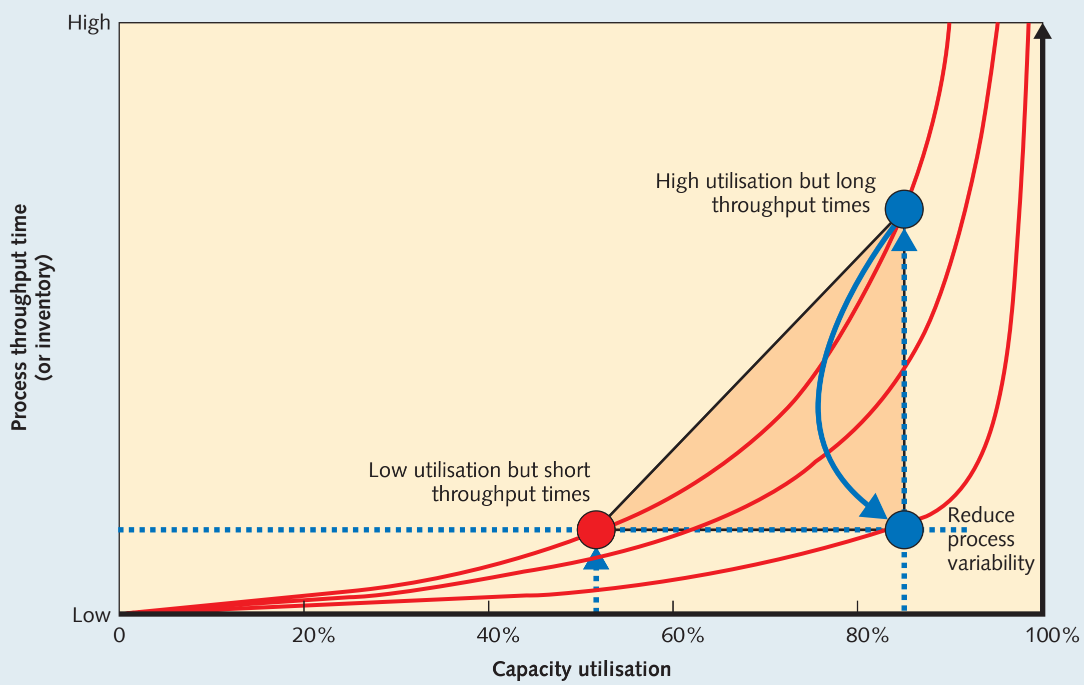

Figure 11.11 Developing lean processes can mean accepting lower utilisation in the short to medium term

towards lean synchronisation, fast throughput time and smooth flow is more important that the high utilisation which can result in build-up of inventory or queues. However, this criticism is not really valid in the long term. Remember the relationship between capacity utilisation and process throughput time (or inventory of products, customers, or information), shown in Figure 11.11. The improvement path envisaged by adopting lean synchronisation is shown as moving from the state that most businesses find themselves in (high utilisation but long throughput times) towards the lean synchronisation ideal (short throughput time). Although, inevitably, this means moving towards a position of lower capacity utilisation, lean synchronisation also stresses a reduction in all types of process variability. As this begins to become reality, the improvement path moves towards the point where throughput time is short and capacity utilisation high. It manages to do this because of the reduction in process variability.

- Not all commentators see lean synchronisation-influenced people-management practices as entirely positive. The JIT approach to people management can be viewed as patronising. It may be, to some extent, less autocratic than some Japanese management practices dating from earlier times. However, it is certainly not in line with some of the job design philosophies which place a high emphasis on contribution and commitment. Even in Japan, the JIT approach is not without its critics. Kamata wrote an autobiographical description of life as an employee at a Toyota plant called Japan in the Passing Lane.[10] His account speaks of 'the inhumanity and the unquestioning adherence' of working under such a system. Some trade union representatives have voiced similar criticisms.

- Any textbook of this type has to segment the ideas and knowledge contained within its subject in such a way that each set of ideas is explained and communicated as clearly as possible. Yet doing this inevitably means imposing artificial boundaries between the various topics. No more so than in the case of lean synchronisation. There are some particularly evangelical proponents of the lean philosophy who object strongly to separating out the whole concept of lean into a separate chapter. The underlying ideas of lean, they say, have now comprehensively replaced those ideas described as 'traditional' at the beginning of this chapter. Rather, lean principles should be the foundation for the whole of operations and process management. Lean principles have something to tell us about everything in the subject from quality management to inventory management, from job design to product design. And they are right of course. Nevertheless, the ideas behind lean synchronisation are both counterintuitive enough and important enough to warrant separate treatment. Also lean in its pure form is not necessarily equally applicable to every situation (refer to the discussion about lean and agile). Hence the inclusion of this chapter that focuses on this topic. Remember though, lean synchronisation is one of those topics (like operations strategy, quality and improvement) that has a particularly strong influence over the whole subject.

SUMMARY CHECKLIST

- ☐ Are the benefits of attempting to achieve lean synchronisation well understood within the business?
- ☐ Notwithstanding that the idea derives from manufacturing operations, have the principles been considered for non-manufacturing processes within the business?
- ☐ Is the extent of waste within operations and processes fully understood?
- ☐ Can flow of items through processes be made more regular?
- ☐ How much inventory (products, customers and information) is building up because of inexact supply?
- ☐ How much waste is caused because of inflexibility in the operation's processes?
- ☐ How much waste is caused because of variability (especially of quality) within the operation's processes?
- ☐ Are capacity utilisation performance measures likely to prove a barrier to achieving lean synchronisation?
- ☐ Does the culture of the organisation encourage the involvement in the improvement process of all people in the organisation?
- ☐ Are the ideas of continuous improvement understood and used in practice?
- ☐ Are the various techniques used to promote lean synchronisation understood and practiced?
- ☐ Is the concept of lean synchronisation applied throughout the supply network?
- ☐ Has the possibility of blending push (such as MRP) and pull (such as lean synchronisation) been considered?

CASE STUDY

Saint Bridget's Hospital[11]

When Denize Ahlgren arrived at St Bridget's, one of the main hospitals in the Götenborg area, she knew that it had gained a reputation for fresh thinking on how healthcare could be organised to give better public care at lower cost to the taxpayer. In fact, that was one of the reasons she had taken the job of its Chief of Administration (COA). In particular, Denize had read about St Botvid's 'Quality Care' (QC) initiative. 'Yes, QC is obviously important', explained Dr Pär Solberg who, in addition to his clinical duties, also headed the QC initiative, 'but don't think that it is only about "quality". We don't just throw money at improving the quality of care; we also want to improve efficiency. Any money saved by improving efficiency can then be invested in improving clinical outcomes.'

'It all started with quality'

Although run by a private company, St Bridget's is little different from any other Swedish hospital. To its patients treatment is free, after a minimal charge that is universal in Sweden. St Bridget's gets virtually all its revenue from the government. However, in terms of how it organises itself, it is at the forefront of implementing ideas that are more common in private business. 'It all started with our efforts a few years ago to be systematic in how we measured quality', said Pär Solberg. 'We felt that quality must be reported on a systematic and logical basis if it is going to be meaningful. It should also be multi-faceted, and not just focus on one aspect of quality. We measure three aspects, "reported patient experience" (RPE), what the patient thinks about the total experience of receiving treatment, "reported patient outcome" (RPO), how the patient views the effectiveness of the treatment received, and most importantly "reported clinical outcome" (RCO), how the clinicians view the effectiveness of the treatment. Of course these three measures are interconnected. So, RPO eventually depends on the medical outcome (RCO) and how much discomfort and pain the treatment triggers. But it is also influenced by the patient's experience (RPE), for example, how well we keep the patient informed, how empathetic our staff are, and so on.'

'Measuring quality led naturally to continuous improvement'

The hospital's quality measurement processes soon developed into a broader approach to improvement in general. In particular, the idea of continuous improvement began to be discussed. 'Measuring quality led naturally to continuous improvement', explained Pär Solberg. 'Once we had measurable indicators of quality, we could establish targets, and most importantly we could start to think about what was preventing us improving quality. This, in turn, led to an understanding of all the processes that affected quality indicators. It was a shift to seeing the hospital as a whole set of processes that governed a set of flows – flows of patients through their treatment stages, flows of clinical staff, flows of information, flows of pharmaceuticals, flows of equipment, and so on. It was a revolution in our thinking. We started examining these flows and looking at how they impacted on our performance and how we could improve the working methods that we considered significant for the quality indicators that we wanted to influence. That was when we discovered the concept of "lean".'

'Continuous improvement introduced us to lean'

It was at an 'Improving European Healthcare' conference, which Pär and another colleague attended that St Bridget's was first introduced to the idea of 'lean'. 'Continuous improvement introduced us to lean. We were talking to some representatives from the UK's National Health Service Institute, who had been involved in introducing lean principles in UK hospitals. They explained that lean was an improvement approach that improved flow and eliminated waste, which had been used successfully in some hospitals to build on continuous improvement. Lean, they said, as developed by Toyota was about getting the right things to the right place, at the right time, in the right quantities, while minimising waste and being flexible and open to change. It sounded worth following up. However they admitted that not every attempt to introduce lean principles had met with success.'

'It can easily all get political'

Intrigued by the conversation, Pär contacted one of the hospitals in the UK that had been mentioned, and talked to Marie Watson who had been the 'Head of Lean' and had initiated several lean projects. She said that one of the problems she had faced was her Chief Executive's insistence on bringing in several firms of consultants to implement lean ideas. To make matters more confusing, when a new Chief Executive was appointed, he brought in his own preferred consultants in addition to those already operating in the hospital. Marie had not been happy with the change. 'Before the change of executives, we had a very clear way of how we were going to move forward and spread lean throughout the organisation, then we became far less clear. The emphasis shifted to get some quick results. But that wasn't why we were set up. Originally it was about having a positive impact, getting people involved in lean, engaging and empowering them towards continuous improvement; there were things that were measurable but then it changed to "show us some quick results". People were forgetting the cultural side of it. Also it can easily all get "political". The different consultancy teams and the internal lean initiatives, all had their own territories. For example, we [Marie's internal team] were about to start a study of A&E activities, when they were told to keep away from A&E so as not to "step on the toes" of the firm of consultants working there.'

'We're not making cars, people are different'

Pär was determined not to make the same mistakes that Marie's hospital had, and consulted widely before attempting any lean improvements with his colleagues. Some were sceptical: 'we're not making cars, people are different and the processes that we put people through repeatedly are more complicated than the processes that you go through to make a car.' Also, some senior staff were dubious about changes that they perceived to threaten their professional status. Instead of doctors and nurses maintaining separate and defined roles that focused solely on their field of medical expertise, they were encouraged to work (and sit) together in teams. The teams were also made responsible for suggesting process improvements. But most could be converted. One senior clinician, at first claimed that, 'this is all a load of rubbish. There's no point in mapping this process, we all know what happens: the patient goes from there to there and this is the solution and that's what we need to do.' Yet only a few days later, he said, 'I never realised this is what really happens, that won't work now will it, actually this has been great because I never understood, I only saw my bit of it, now I understand all of the process.'

'It works, it makes things better for the patients'

Over time, most (although not quite all) scepticism was overcome, mainly because, in the words of one doctor, 'It works, it makes things better for the patients.' As more parts of the hospital became convinced of the effectiveness of the lean approach, the improvements to patient flow and quality started to accumulate. Some of the first improvements were relatively simple, such as a change of signage (to stop patients getting lost). Another simply involved a roll of yellow tape. Rather than staff wasting precious time looking for equipment such as defibrillators, the yellow tape was used to mark a spot on the floor where the machines were always kept. Another improvement involved using magnetic dots on a progress chart to follow each patient's progress and indicate which beds were free. Some were even simpler, for example discharging patients throughout the day rather than all at the same time, so that they can easily find a taxi. Other improvements involved more analysis, such as reducing the levels of stock being held (e.g., 25,000 pairs of surgical gloves from 500 different suppliers). Some involved a complete change in assumptions, such as the effectiveness of the medical records department. 'It was amazing. We just exploded the myth that when you didn't get case notes in a clinical area it was medical records fault. But it never was. Medics had notes in their cars, they had them at home, we had a thousand notes in the secretaries' offices, there were notes in wards, drawers and cupboards, they were all over the place. And we wondered why we couldn't get case notes! Two people walked 7 miles a day to go and find case notes, for case!' (Pär Solberg)

'We need to go to the next level'

Denize Ahlgren was understandably impressed by the improvements that Pär had outlined to her; however, Pär was surprisingly downbeat about the future. 'OK, I admit that we have had some impressive gains from continuous improvement and latterly from the adoption of lean principles. I am especially impressed with Toyota's concept of the seven types of waste (for more details see 'Example – Toyota' earlier in this chapter). It is both a conceptually powerful and a very practical idea for identifying where we could improve. Also the staff like it. But it's all getting like a box-ticking exercise. Looking for waste is not exactly an exciting or radical idea. The more that I study how lean got going in Toyota and other manufacturing plants, the more I see that we haven't really embraced the whole philosophy. Yet, at the same time, I'm not totally convinced that we can. Perhaps some of the doubters were right; a hospital isn't a car plant, and we can apply only some lean ideas.'

Ironically, as Pär was having doubts some of colleagues were straining to do more. One clinician, in particular, Fredrik Olsen, chief physician at St Botvid's lower back pain clinic, thought that his clinic could benefit from a more radical approach. 'We need to go to the next level. The whole of Toyota's philosophy is concerned with smooth synchronous flow, yet we haven't fully got our heads round that here. I know that we are reluctant to talk about "inventories" of patients, but that is exactly what waiting rooms are. They are "stocks" of people, and we use them in exactly the same way as pre-lean manufacturers did – to buffer against short-term mismatches between supply and demand. What we should be doing is tackling the root causes of the mismatch. Waiting rooms are stopping us from moving towards smooth, value-added, flow for our patients.'

Fredrik went on to make what Denize thought was an interesting, but radical, proposal. He proposed scrapping the current waiting room for the lower back pain clinic and replacing it with two extra consulting rooms to add to the two existing consulting rooms. Patients would be given appointments for specific times rather than being asked to arrive 'on the hour' (effectively in batches) as at present. A nurse would take the patents' details and perform some preliminary tests, after which they would call in the specialist physician. Staffing levels during clinic times would be controlled by a nurse who would also monitor patient arrival, direct them to consulting rooms and arrange any follow-up appointments (for MRI scans for example).

Denize was not sure about Fredrik's proposal. 'It seems as though it might be a step too far. Patients expect to wait until a doctor can see them, so I'm not sure what benefits would result from the proposal. And what is the point of equipping two new consulting rooms if they are not going to be fully utilised?'

QUESTIONS

1 What benefits did St Bridget's get from adopting first a continuous improvement, then a lean, approach?

2 Do you think that Pär Solberg is right in thinking that there is a limit to how far a hospital can go in adopting lean ideas?

3 On the St Bridget's website there are several references to its 'Quality Care' programme, but none to its lean initiatives, even though lean is regarded as important by most clinicians and administrators in the hospital. Why do you think this might be?

4 Denize cannot see the benefits of Fredrik's proposal. What do you think they might be?

5 Are any benefits of scrapping the waiting room in the clinic worth the underutilisation of the four consulting rooms that Fredrik envisages?

APPLYING THE PRINCIPLES

Some of these exercises can be answered by reading the chapter. Others will require some general knowledge of business activity and some might require an element of investigation. **Model answers for the first two questions are available on the student companion website.** For model answers to the other questions in this section, please ask your tutor/lecturer.

1 Re-examine the description of the Toyota production system at the beginning of the chapter.

(a) List all the different techniques and practices that Toyota adopts. Which of these would you call just-in-time philosophies and which are just-in-time techniques?

(b) How are operations objectives (quality, speed, dependability, flexibility, cost) influenced by the practices which Toyota adopts?

2 Consider this record of an ordinary flight. '*Breakfast was a little rushed but left the house at 6.15. Had to return a few minutes later, forgot my passport. Managed to find it and leave (again) by 6.30. Arrived at the airport 7.00, dropped Angela off with bags at terminal and went to the long-term car park. Eventually found a parking space after 10 minutes. Waited 8 minutes for the courtesy bus. Six-minute journey back to the terminal; we started queuing at the check-in counters by 7.24. Twenty minute wait. Eventually get to*

check-in and find that we have been allocated seats at different ends of the plane. Staff helpful but takes 8 minutes to sort it out. Wait in queue for security checks for 10 minutes. Security decide I look suspicious and search bags for 3 minutes. Waiting in lounge by 8.05. Spend 1 hour and 5 minutes in lounge reading computer magazine and looking at small plastic souvenirs. Hurrah, flight is called 9.10, takes 2 minutes to rush to the gate and queue for further 5 minutes at gate. Through the gate and on to air bridge that is a continuous queue going onto plane; takes 4 minutes but finally in seats by 9.21. Wait for plane to fill up with other passengers for 14 minutes. Plane starts to taxi to runway at 9.35. Plane queues to take-off for 10 minutes. Plane takes off at 9.45. Smooth flight to Amsterdam, 55 minutes. Stacked in queue of planes waiting to land for 10 minutes. Touch down at Schiphol Airport 10.50. Taxi to terminal and wait 15 minutes to disembark. Disembark at 11.05 and walk to luggage collection (calling at lavatory on way); arrive luggage collection 11.15. Wait for luggage 8 minutes. Through customs (not searched by Netherlands security who decide I look trustworthy) and to taxi rank by 11.26. Wait for taxi 4 minutes. In to taxi by 11.30; 30 minutes ride into Amsterdam. Arrive at hotel 12.00.'

(a) Analyse the journey in terms of value-added time (actually going somewhere) and non-value-added time (the time spent queuing, etc.).

(b) Visit the websites of two or three airlines and examine their business class and first class services to look for ideas that reduce the non-value-added time for customers who are willing to pay the premium.

(c) Next time you go on a journey, time each part of the journey and perform a similar analysis.

3 An insurance underwriting process consists of the following separate stages.

What is the 'value-added' percentage for the process? (Hint – use Little's Law to work out how long applications have to wait at each stage before they are processed. Little's Law is covered in Chapter 6.)

Stage	*Processing time per application (minutes)*	*Average work-in-progress before the stage*
Data entry	30	250
Retrieve client details	5	1,500
Risk assessment	18	300
Inspection	15	150
Policy assessment	20	100
Dispatch proposal	10	100

4 Examine the value-added versus non-value-added times for some other services. For example,

(a) Handing-in an assignment for marking if you are currently studying for a qualification. (What is the typical elapsed time between handing the assignment in and receiving it back with comments?) How much of this elapsed time do you think is value-added time?

(b) Posting a letter (the elapsed time is between posting the letter in the box and it being delivered to the recipient).

5 Go back to Chapter 10 and re-read the description of the Theory of Constraints (ToC). Now consider the similarities and differences between ToC and lean synchronisation in terms of their overall objectives, measures of effectiveness, improvement method and implementation.

Notes on chapter

1 Sources include: Spears, S. and Bowen, H.K. (1999) 'Decoding the DNA of the Toyota Production System', *Harvard Business Review, October, 96-106.*

2 Sources include: Interview with Edward Kay, Tom Dyson and Olly Willans of Torchbox; Torchbox website, http://www.torchbox.com/; we are grateful to everyone at Torchbox for their help and allowing us access to their operation.

3 Lee, D.C. (1987) 'Set-up time reduction: making JIT work' *in* Voss, C.A. (ed), *Just-in-time Manufacture,* IFS/Springer-Verlag.

4 Example written and supplied by Janina Aarts and Mattia Bianchi, Department of Management and Organization, Stockholm School of Economics

5 Spears, S. and Bowen, H.K. op.cit

6 Harrison, A. (1992) *Just-in-time Manufacturing in Perspective,* Prentice Hall.

7 Sources include, *McKinsey Quarterly* (2014) 'When Toyota met e-commerce: lean at Amazon', no. 2; Liker, J. (2004) *The Toyota Way,* McGraw-Hill; Rosenthal, M. (2002) 'The essence of Jidoka', *SME Lean Directions Newsletter.*

8 This great metaphor seems to have originated from the consultancy '2think', http://www.2think.biz/index.htm

9 Kruse, G. (2002) 'IT enabled lean agility', *Control,* November.

10 Kamata, S. (1983) *Japan in the Passing Lane,* Allen and Unwin.

11 This case is based on the work of several real hospitals in Scandinavia and the rest of the world that have used the concepts of lean operations to improve their performance. However, all names and places are fictional and no connection is intended to any specific hospital.

TAKING IT FURTHER

Bicheno, J. and Holweg, M. (2016) *The Lean Toolbox: The Essential Guide to Lean Transformation (5th edn), PICSIE Books.* *A practical guide from two of the European authorities on all matters lean.*

Holweg, M. (2007) *'The genealogy of lean production', Journal of Operations Management, vol. 25, 420–437.* *An excellent overview of how lean ideas developed.*

Mann, D. (2010) *Creating a Lean Culture (2nd edn), Productivity Press.* *Treats the soft side of lean.*

Modig, N. and Ahlstrom, P. (2012) *This is Lean: Resolving the Efficiency Paradox, Rheologica Publishing.* *This book provides a very practical guide to what lean is and its application in a variety of sectors. Not only does this book demonstrate a clear understanding of how the various aspects of lean come together, it does this in a very readable way.*

Womack, J.P., Jones, D.T. and Roos, D. (1990) *The Machine that Changed the World, Rawson Associates.* *Arguably the most influential book on operations management practice of the last fifty years. Firmly rooted in the automotive sector but did much to establish JIT.*

Womack, J. P. and Jones, D. T. (2003) *Lean Thinking: Banish Waste and Create Wealth in Your Corporation, Free Press.* *Some of the lessons from 'The Machine that Changed the World' but applied in a broader context.*

12 Improvement

Introduction

Improvement means to make something better. And all operations, no matter how well managed, are capable of improvement. At one time, operations managers were expected simply to ‘run the operation’, ‘keep the show on the road’ and ‘maintain current performance’. No longer is this true. In fact, in recent years the emphasis has shifted markedly towards making improvement one of the main responsibilities of operations managers. And although the whole of this book is focused on improving the performance of individual processes, operations and whole supply networks, there are some issues that relate to the activity of improving itself. In any operation, whatever is improved, and however it is done, the overall direction and approach to improvement needs to be addressed. Figure 12.1 shows the position of the ideas described in this chapter in the general model of operations management.

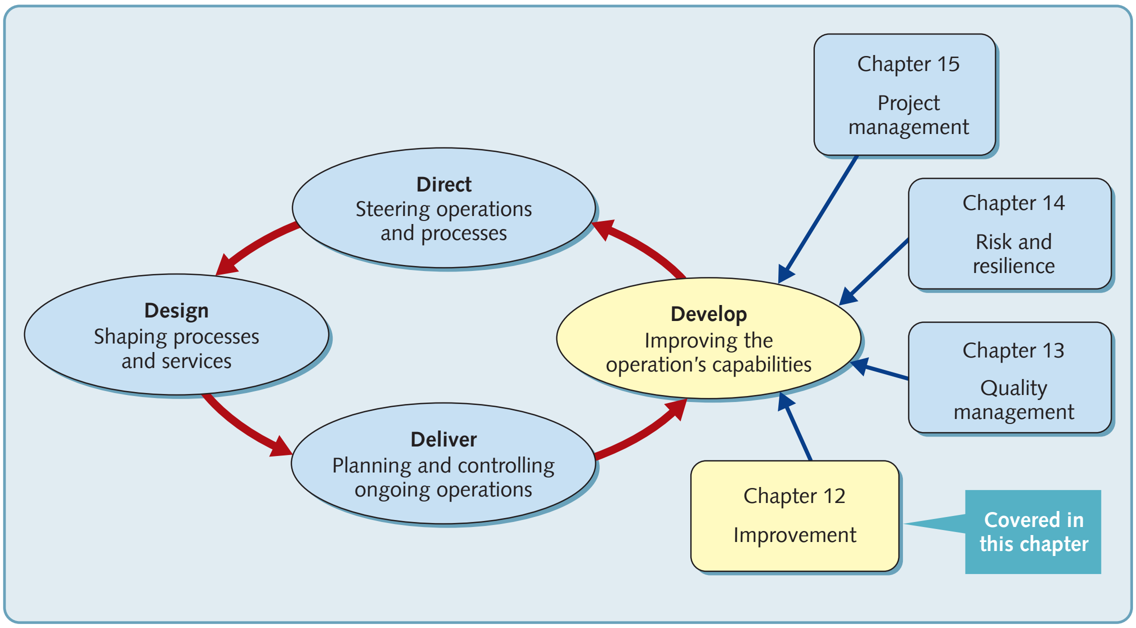

Figure 12.1 Improvement is the activity of closing the gap between the current and the desired performance of an operation or process.

EXECUTIVE SUMMARY

Why is improvement so important?
↓
What is the gap between current and required performance?
↓
What is the most appropriate improvement path?
↓
What techniques are being used for improvement?
↓
Does improvement always stick?

Why is improvement so important?

Improvement is the activity of closing the gap between the current and the desired performance of an operation or process. It is increasingly seen as the ultimate objective for all operations and process management activity. Furthermore, almost all popular operations initiatives in recent years, such as total quality management, lean operations, business process re-engineering and Six Sigma, have all focused on performance improvement. It involves assessing the gaps between current and required performance, balancing the use of continuous improvement and breakthrough improvement, adopting appropriate improvement techniques and attempting to ensure that the momentum of improvement does not fade over time.

What is the gap between current and required performance?

Assessing the gap between actual and desired performance is the starting point for most improvement. This requires two sets of activities: first, assessing the operation's and each process's current performance; and second, deciding on an appropriate level of target performance. The first activity will depend on how performance is measured within the operation. This involves deciding what aspects of performance to measure, which are the most important aspects of performance and what detailed measures should be used to assess each factor. The balanced score card approach is an approach to performance measurement that is currently influential in many organisations. Setting targets for performance can be done in a number of ways. These include historically based targets, strategic targets that reflect strategic objectives, external performance targets that relate to external and/or competitor operations and absolute performance targets based on the theoretical upper limit of performance. Benchmarking is an important input to this part of performance improvement.

What is the most appropriate improvement path?

Two improvement paths represent different philosophies of improvement, although both may be appropriate at different times. They are breakthrough improvement and continuous improvement. Breakthrough improvement focuses on major and dramatic changes that are intended to result in dramatic increases in performance. The business process re-engineering approach is typical of breakthrough improvement. Continuous improvement focuses on small but never-ending improvements that become part of normal operations life. Its objective is to make improvement part of the culture of the organisation. Often continuous improvement involves the use of multi-stage improvement cycles for regular problem solving. The Six Sigma approach to improvement brings many existing ideas together and can be seen as a combination of continuous and breakthrough improvement.

What techniques should be used to facilitate improvement?

Almost all techniques in operations management contribute directly or indirectly to the performance improvement. However, some more general techniques have become popularly associated with improvement. These include scatter diagrams (correlation), cause–effect diagrams, Pareto analysis and why–why analysis.

How can improvement be made to stick?

One of the biggest problems in improvement is to preserve improvement momentum over time. One factor that inhibits improvement becoming accepted as a regular part of operations activity is the emphasis on the fashionability of each new improvement approach. Most new improvement ideas contain some worthwhile elements but none will provide the ultimate answer. There must be some overall management of the improvement process that can absorb the best of each new idea. And, although authorities differ to some extent, most emphasise the importance of an improvement strategy, top-management support and training.

DIAGNOSTIC QUESTION

Why is improvement so important?

Why is operations improvement so important? Well, who doesn't want to get better? And businesses are (or should be) just the same as people – they generally want to get better. Not just for the sake of their own excellence, although that may be one factor, but mainly because improving operations performance has such an impact on what any organisation is there to do. Emergency services want to reach distressed people faster and treat them better because by doing so they are fulfilling their role more effectively. Package delivery businesses want to deliver more reliably, at lower cost and reducing emissions because it means happier customers, higher profits and less pollution. Development charities want to target their aid campaign for improvement in human conditions as wisely and efficiently as possible because more money will find its way to beneficiaries, rather than be wasted or consumed in administration. It is not surprising then that operations management has shifted towards emphasising improvement. Operations managers are judged not only on how they meet their ongoing responsibilities of producing products and services to acceptable levels of quality, speed, dependability, flexibility and cost, but also on how they improve the performance of the operations function overall.

Improvement comes from closing the gap between what you are and what you want to be. Or, in a specifically operations context, it comes from closing the gap between current and desired performance. Performance improvement is the ultimate objective of operations and process management. It has also become the subject of innumerable ideas that have been put forward as particular effective methods of ensuring improvement. These include many that are described in this book, for example, total quality management (TQM), 'lean' operations, business process re-engineering (BPR), Six Sigma, and so on. All of these, and other, ideas have something to contribute. What is important is that all managers develop an understanding of the underlying elements of improvement. The following two examples illustrate many of these.

OPERATIONS PRINCIPLE

Performance improvement is the ultimate objective of operations and process management.

Why the focus on improvement?

Various reasons have been suggested to explain the shift towards a focus on improvement in professional operations managers' activities.

- There is a perceived increase in the intensity of competitive pressures (or 'value for money in not-for-profit or public sector operations). In fact, economists argue about whether markets are really getting more competitive. As far as improvement is concerned it doesn't matter; there is a *perception* of increased competitive pressure, and certainly the owners of operations (shareholders or governments) are less likely to tolerate poor returns or value for money.
- The nature of world trade is changing. Economies such as China, India and Brazil are emerging as both producers and consumers of products and services. This has had a number of effects that have impacted more developed economies. It has introduced cost pressures in countries with relatively expensive labour and infrastructure costs; it has introduced new challenges for global companies, such as managing complex supply chains; and it has accelerated demand for resources (materials, food, energy), pushing up (or destabilising) prices for these commodities.

- New technology has both introduced opportunities to improve operations practice and disrupt existing markets. Look at how operations in the music business have had to adapt their working practices to downloading and music streaming.
- The interest in operations improvement has resulted in the development of many new ideas and approaches to improving operations that have, in turn, focused attention on improvement. The more ways there are to improve operations, the more operations will be improved.

EXAMPLE

Building improvement at Prêt A Manger[1]

Prêt A Manger, is a chain of selling sandwiches and similar lunch products, founded in London in the late 1980s by college friends Julian Metcalfe and Sinclair Beecham. They now have shops through the UK, and increasingly in other large cities such as Paris, New York, Boston, Hong Kong and Tokyo. Although Prêt struggled to make money at first, they eventually developed a range of products and an operating model that worked. And it has continued to work. Described by the press as having 'revolutionised the concept of sandwich-making and eating', their revenues per store can be 50 per cent higher than rivals. It is a success that is a result of the company's unrelenting focus on improvement in three areas.

The first is the quality, taste, affordability and overall attractiveness of the food it sells. Julian Metcalfe, one of the founders, is a self-confessed 'obsessive about improvement'. '*The difference between mediocre and great* [is] *extremely elusive and some of our sandwiches are bloody great, they really work. It takes years of relationships with the suppliers to get the right cheese, to get the right seasoning mix in the mayo.*' The company constantly reinvents its products, even those that remain popular; for example, Prêt's pickle recipe has been revised 15 times, its chocolate brownie 36 times and its carrot cake 50 times. They avoid the chemicals and preservatives common in most 'fast' food. '*Many food retailers focus on extending the shelf-life of their food, but that's of no interest to us. We sell food that can't be beaten for freshness. At the end of the day, we give whatever we haven't sold to charity to help feed those who would otherwise go hungry.*' The company also runs a training scheme for homeless people and ex-offenders, and said that more than 70 people have been 'taken off the streets' as a result.

The second focus of Prêt's improvement efforts is maintaining and enriching its (already outstanding) standards of customer service. '*We'd like to think we react to our customers' feelings (the good, the bad, the ugly) with haste and absolute sincerity*', they said., *Prêt customers have the right to be heard. Do call or email. Our UK Managing Director is available if you would like to discuss Prêt with him. Alternatively, our CEO hasn't got much to do; hassle him!*' Prêt A Manger shops have their own kitchen where fresh ingredients are delivered every morning, with food prepared throughout the day. The team members serving on the tills at lunchtime will have been making sandwiches in the kitchen that morning. '*We are determined never to forget that our hardworking people make all the difference. They are our heart and soul. When they care, our business is sound. If they cease to care, our business goes down the drain. In a retail sector where high staff turnover is normal, we're pleased to say our people are much more likely to stay around! We work hard at building great teams. We take our reward schemes and career opportunities very seriously. We don't work nights (generally), we wear jeans, we party!*' Potential employees are often asked to work for a day in a Prêt shop; after which their fellow workers vote on whether or not they should be hired. Also, every Prêt manager is required to spend four days a year on the shop floor. Customer feedback is regarded as being particularly important at Prêt. Examining customers' comments for improvement ideas is a key part of weekly management meetings, and of the daily team briefs in each shop. Moreover, staff at Prêt are rewarded in cash for being nice to customers. They collect bonuses for delivering outstanding customer service. Every week, each Prêt outlet is visited by a 'mystery shopper' who scores the shop on such performance measures as speed of service, product availability and cleanliness. In addition, the mystery shopper rates the 'engagement level' of the staff; questions include, 'did servers connect with eye contact, a smile and some polite remarks?' Assessors score out of 50. If the store gets 43 points or more every team member receives an extra payment for every hour worked; and if an individual is mentioned by the mystery shopper for providing outstanding service, they get an extra payment. The reward is for service rather than sales. '*The emphasis on jollity and friendliness has been a winner*', said James Murphy of the Future Foundation, a management consultant. '*In the highly competitive sandwich market, that's been a big contributor to their success.*'

The third area of improvement concerns how new store outlets are added to its network, so that customers continue to find their stores 'the convenient option'. Clive Schlee, Prêt's chief executive, said that the group's expansion plans remain flexible. '*We've always managed expansion quite carefully*', he said. '*We don't have to put in big fixed investments.*' Nor do Prêt use franchising to expand. '*When you begin to franchise out,*' said Tracy Gingell, a manager at New York's Broad Street shop, '*what's to stop a store manager from buying the five-dollar case of chicken instead of the $35 case that's organic. No one's going to know the difference. That's why we don't franchise.*' But the network expansion aspect of Prêt's improvement philosophy has not always been a success. In the early 2000s, after an over-hasty expansion, quality suffered, and their recently opened Japanese stores were closed.

EXAMPLE

Improving Sonae Corporation's retail operations

(Adapted, with permission, from an original case by Professors Rui Soucasaux Sousa and Sofia Salgado Pinto, Católica Porto Business School, Portugal)

The retail industry may not seem to be the most likely setting for using improvement approaches more usually associated with manufacturing, but Sonae Corporation's Continente supermarkets' 22,500 employees in its 170 stores and two distribution centres have demonstrated that any type of business can benefit. The operations improvement programme was originally a response to Portuguese labour laws requiring a minimum of 35 hours of training per year, per worker. Jaime Maia, Sonae's human resource director was keen that the training should be 'on-the-job' and asked a consulting firm, The Kaizen Institute (KI) to help, first by observing daily operations at several retail stores. The results were surprising. KI uncovered a significant amount of waste (or 'muda' in lean terminology, see Chapter 11). However, when Jaime Maia and KI presented their ideas to the chief operations officer (COO), their photographs showing examples of waste caused some discomfort; after all, they were the most successful retailer in Portugal. Yet the program went ahead, despite several managers arguing that lean principles would not work in retail.

Within a store, 'back-office' processes include unloading trucks, routing goods to sales areas or the warehouses, cleaning, shelf replenishment and store decoration. 'Front office' processes include sales areas and their supporting checkout and customer service areas. Key operational goals include the efficient use of space, increasing sales per square metre of store space and customer satisfaction. The first stage of the programme focused on 'goods reception' and 'shelf replenishment'. Simple lean tools such as 5S and Visual Management (see Chapter 15) were used, as was the idea of 'Gemba', or working out improvements in the workplace. After a seminar, store managers returned to work with their teams to develop training based on a 'lean manual', define an action plan, identify problem areas, and select the lean tools to be applied. As Jaime Maia explained, '*Improvements were suggested by store managers, top-down. But those ideas were immediately enriched and put into action by the teams in the stores, bringing about further improvements in a continuous fashion.*'

After just one year there had been an 'explosion of creativity' in the stores. Productivity had increased; inventory and stock-outs were reduced, and customer satisfaction increased. As Jaime Maia put it '[continuous improvement] *stimulates a good attitude and a constant sense of critique*'. A typical improvement project concerned the company's shelf-replenishing policy. Initially, stock was continuously replenished as sales took place during the day. However, this meant that product movements were constrained by customer flows and the need to keep the store clean and tidy. So a new method was tested. '*The store is fully loaded before the morning opening. From then on, we just need to perform minimal stock maintenance during the day. There is a time of the day at which a shelf may appear to be quite empty. However, typically, there is no need to replenish the shelf, but simply bring the products from the back to the front of the shelf, or from the upper shelves to the eye level shelves*', explained Nuno Almeida, regional operations manager. Only a few fast-moving goods needed to be replenished during the day with the store open.

With the success of the programme it was expanded by involving all employees in the program. A formal steering group was created with monthly general meetings and a videoconference meeting every two weeks to assess progress. Substantial improvements in performance continued, but progress was not uniform. Stores, ranged from 87 per cent implementation, down to 37 per cent implementation. So, the steering group decided to place more emphasis on benchmarking and learning.

And the future? '*One challenge*', said Jaime Maia '*is to sustain the motivation for the programme across the organisation, after years of continuous successes'*. He also felt the program was reaching a new turning point and needed to be reinvented. Until now, lean principles had been applied mainly to materials flows and workplace organisation. Could lean principles be extended to customer flows?

What do they have in common?

The improvement initiatives at these two operations, and the way they managed them, is typical of improvement projects. Both measured performance, implicitly placing information gathering at the centre of their improvement initiative. Both had a view of improvement targets that related directly to strategic objectives. Both made efforts to collect information that would allow decisions based on evidence rather than opinion. Prêt A Manger also made extensive use of mystery shoppers to check that their standards were being maintained. Sonae used kaizen (see later) and several elements of lean (see Chapter 11) to focus on each individual site's improvement. They both had to foster an environment that allowed all staff to contribute to its improvement, and both came to view improvement not as a 'one-off', but rather as the start of a never-ending cycle of improvement. Most importantly, both had to decide how to organise the whole improvement initiative. Different organisations with different objectives may choose to implement improvement initiatives in a different way. But all will face a similar set of issues to these two operations, even if they choose to make different decisions.

The Red Queen effect

The scientist Leigh Van Valen was looking to describe a discovery that he had made while studying marine fossils. He had established that, no matter how long a family of animals had already existed, the probability that the family will become extinct is unaffected. In other words, the struggle for survival never gets easier. However well a species fits with its environment, it can never relax. The analogy that Van Valen drew came from 'Alice's adventures through the looking glass', by Lewis Carroll.[2] In the book, Alice had encountered living chess pieces and, in particular, the 'Red Queen'. '*Well, in our country'*, said Alice, still panting a little*, 'you'd generally get to somewhere else – if you ran very fast for a long time, as we've been doing*'. '*A slow sort of* *country!'* said the Queen. *'Now, here, you see, it takes all the running you can do, to keep in the same place. If you want to get somewhere else, you must run at least twice as fast as that!'*

In many respects this is like business. Improvements and innovations may be imitated or countered by competitors. For example, in many competitive sectors such as consumer electronics or the automotive sector, the quality of most firms' products is significantly better than it was two decades ago. This reflects the improvement in those firm's operations processes. Yet their relative competitive position has in many cases not changed. Those firms that have improved their competitive position have improved their operations performance more than competitors. Where improvement has simply matched competitors, survival has been the main benefit. The implications for operations improvement are clear. It is even more important, especially when competitors are actively improving their operations.

DIAGNOSTIC QUESTION

What is the gap between current and required performance?

The gap between how an operation or process is currently performing, and how it wishes to perform, is the key driver of any improvement initiative. The wider the gap, the more importance is likely to be given to improvement. But, in order to harness the gap as a driver of improvement, it must be addressed in some detail, both in terms of exactly what is failing to meet targets, and by how much. Answering these questions depends on the operation's ability to do three things: assess its current performance; derive a set of target levels of performance that the organisation can subscribe to; and compare current against target performance in a systematic and graphic manner that demonstrates to everyone the need for improvement.

Assessing current performance – performance measurement

Some kind of *performance measurement* is a prerequisite for judging whether an operation is good, bad or indifferent, although this is not the only reason for investing in effective performance measurement. Without one, it would be impossible to exert any control over an operation on an ongoing basis. However, a performance measurement system that gives no help to ongoing improvement is only partially effective. Performance measurement, as we are treating it here, concerns three generic issues.

OPERATIONS PRINCIPLE

Performance measurement is a prerequisite for the assessment of operations performance.

1. Which factors to include as performance measures?
2. Which are the most important performance measures?
3. What detailed measures to use?

Which factors to include as performance measures?

An obvious starting point for deciding which performance measures to adopt is to use the five generic performance objectives: quality, speed, dependability, flexibility and cost. These can be broken down into more detailed measures, or they can be aggregated into 'composite' measures, such as 'customer satisfaction', 'overall service level', or 'operations agility'. These composite measures may be further aggregated by using measures such as 'achieve market objectives', 'achieve financial objectives', 'achieve operations objectives' or even 'achieve overall strategic objectives'. The more aggregated performance measures have greater strategic relevance in so

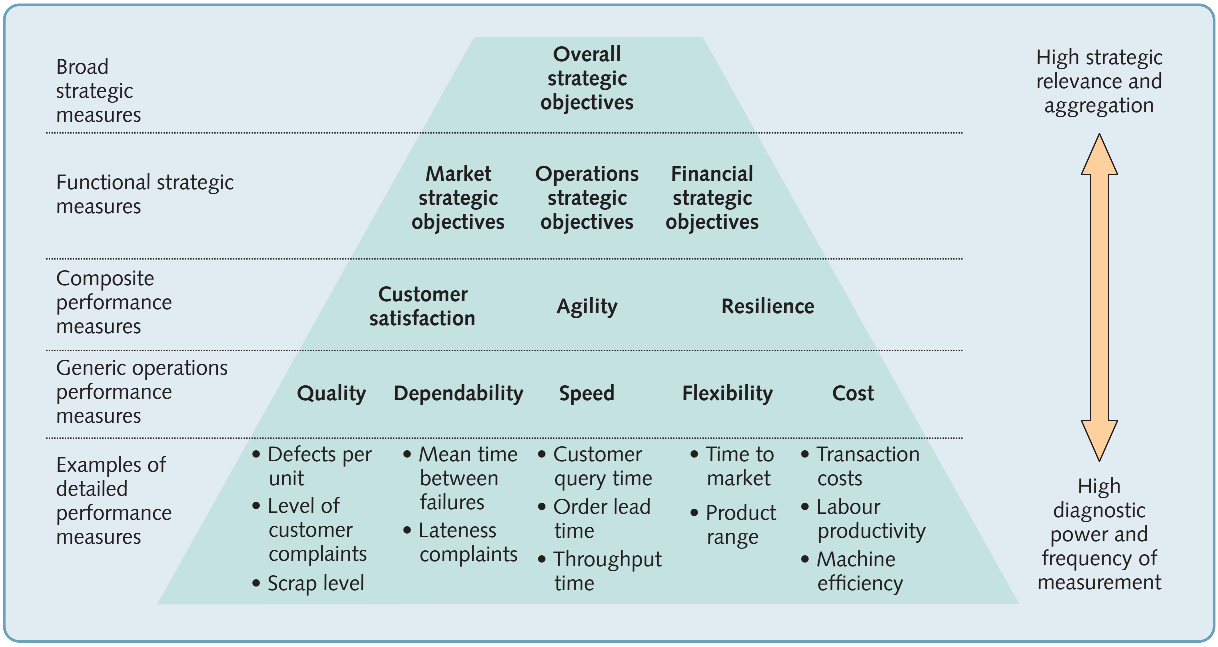

Figure 12.2 **Performance measures can involve different levels of aggregation**

much as they help to draw a picture of the overall performance of the business, although by doing so they necessarily include many influences outside those that operations performance improvement would normally address. The more detailed performance measures are usually monitored more closely and more often; and although they provide a limited view of an operation's performance, they do provide a more descriptive and complete picture of what should be and what is happening within the operation. In practice, most organisations will choose to use performance targets from throughout the range. This idea is illustrated in Figure 12.2.

Choosing the important performance measures

One of the problems of devising a useful performance measurement system is trying to achieve some balance between having a few key measures on the one hand (straightforward and simple, but may not reflect the full range of organisational objectives), or, on the other hand, having many detailed measures (complex and difficult to manage, but capable of conveying many nuances of performance). Broadly, a compromise is reached by making sure that there is a clear link between the operation's overall strategy, the most important (or 'key') performance indicators (KPIs) that reflect strategic objectives, and the bundle of detailed measures that are used to 'flesh out' each key performance indicator. Obviously, unless strategy is well defined then it is difficult to 'target' a narrow range of key performance indicators.

OPERATIONS PRINCIPLE

Without strategic clarity, key performance indicators cannot be appropriately targeted.

What detailed measures to use?

The five performance objectives – quality, speed, dependability, flexibility and cost – are really composites of many smaller measures. For example, an operation's cost is derived from many factors, which could include the purchasing efficiency of the operation, the efficiency with which it converts materials, the productivity of its staff, the ratio of direct to indirect staff, and so on. All of these measures individually give a partial view of the operation's cost performance, and many of them overlap in terms of the information they include. However, each of them does give a perspective on the cost performance of an operation that could be useful either to identify areas for improvement, or to monitor the extent of improvement. If an organisation regards its 'cost' performance as unsatisfactory, disaggregating it into 'purchasing efficiency', 'operations efficiency',

Table 12.1 Some typical partial measure of performance

Performance objective	*Some typical measures*
Quality	Number of defects per unit Level of customer complaints Scrap level Warranty claims Mean time between failures Customer satisfaction score
Speed	Customer query time Order lead time Frequency of delivery Actual *versus* theoretical throughput time Cycle time
Dependability	Percentage of orders delivered late Average lateness of orders Proportion of products in stock Mean deviation from promised arrival Schedule adherence
Flexibility	Time needed to develop new products/services Range of products/services Machine changeover time Average batch size Time to increase activity rate Average capacity/maximum capacity Time to change schedules
Cost	Minimum delivery time/average delivery time Variance against budget Utilisation of resources Labour productivity Added value Efficiency Cost per operation hour

'staff productivity', and so on, might explain the root cause of the poor performance. Table 12.1 shows some of the partial measures which can be used to judge an operation's performance.

Performance measurement and performance management

A distinction that can cause some confusion is that between performance *measurement* and performance *management.* Performance measurement is the attempt to assess of how well (or not) an operation is performing. Performance management is all the actions taken to affect performance in some way. So performance management is what operations managers do to achieve the measures of performance that they are aiming for. And if this seems like a subtle distinction, it is. In fact, measurement and management are not separable.[3] The two things follow one another in an iterative process. Measures of performance reveal where improvement is necessary, and therefore what may need to be done, which is then assessed through measurement. Performance management is a broader issue than performance measurement. It is a philosophy that is supported by performance measurement.

The balanced scorecard approach

'The balanced scorecard retains traditional financial measures. But financial measures tell the story of past events, an adequate story for industrial age companies for which investments in long-term capabilities are customer relationships were not critical for success. These financial

measures are inadequate, however, for guiding and evaluating the journey that information age companies must make to create future value through investment in customers, suppliers, employees, processes, technology, and innovation.'[4]

Generally, operations performance measures have been broadening in their scope. It is now generally accepted that the scope of measurement should, at some level, include external as well as internal, long-term as well as short-term, and 'soft' as well as 'hard' measures. The best-known manifestation of this trend is the 'balanced scorecard' approach taken by Kaplan and Norton.[5] As well as including financial measures of performance, in the same way as traditional performance measurement systems, the balanced scorecard approach also attempts to provide the important information that is required to allow the overall strategy of an organisation to be reflected adequately in specific performance measures. In addition to financial measures of performance, it also includes more operational measures of customer satisfaction, internal processes, innovation and other improvement activities. In doing so, it measures the factors behind financial performance which are seen as the key drivers of future financial success. In particular, it is argued, that a balanced range of measures enables managers to address the following questions. (See Figure 12.3.)

- How do we look to our shareholders (financial perspective)?
- What must we excel at (internal process perspective)?
- How do our customers see us (the customer perspective)?
- How can we continue to improve and build capabilities (the learning and growth perspective)?

The balanced scorecard attempts to bring together the elements that reflect a business's strategic position, including product or service quality measures, product and service development times, customer complaints, labour productivity, and so on. At the same time it attempts to avoid performance reporting becoming unwieldy by restricting the number of measures and focusing especially on those seen to be essential. The advantages of the approach are that it presents an overall picture of the organisation's performance in a single report, and by being comprehensive in the measures of performance it uses, encourages companies to take decisions in the interests of the whole organisation, rather than sub-optimising around narrow measures. Developing a balanced scorecard is a complex process and is now the subject of considerable debate. One of the key questions that has to be considered is how specific measures of

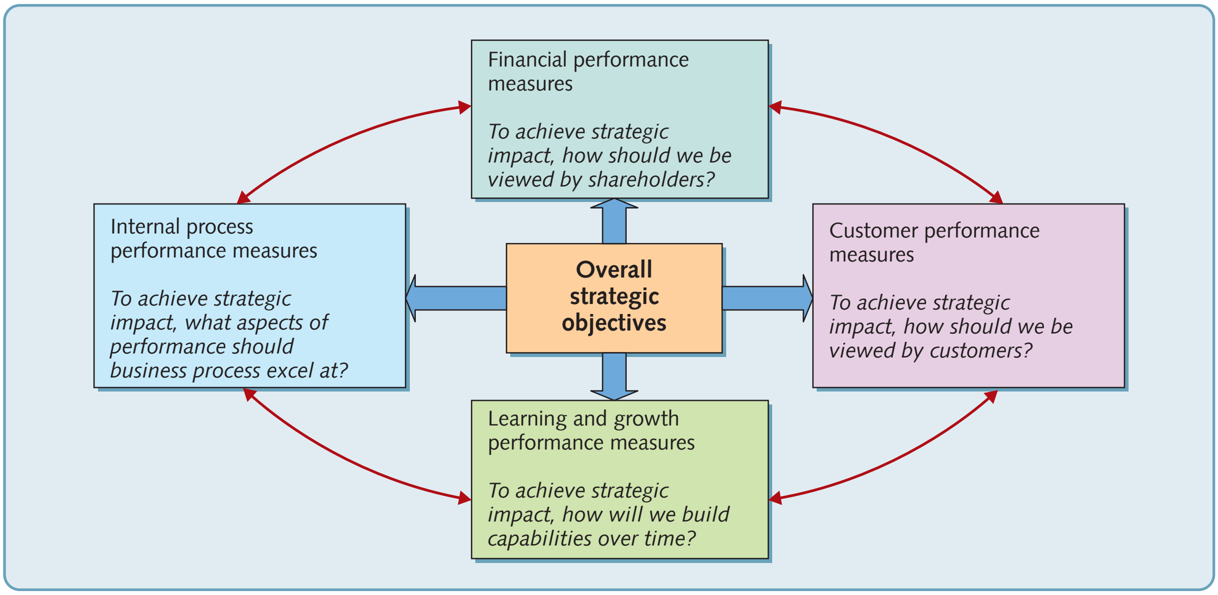

Figure 12.3 The measures used in the balanced scorecard

performance should be designed. Inadequately designed performance measures can result in dysfunctional behaviour, so teams of managers are often used to develop a scorecard that reflects their organisation's specific needs.

Setting target performance

A performance measure means relatively little until it is compared against some kind of target. Knowing that only one document in five hundred is sent out to customers containing an error tells us relatively little, unless we know whether this is better or worse than we were achieving previously, and whether it is better or worse than other similar operations (especially competitors) are achieving. Setting performance targets transforms performance measures into performance 'judgements'. Several approaches to setting targets can be used, including the following:

- *Historically based targets* compare current against previous performance.
- *Strategic targets* are set to reflect the level of performance that is regarded as appropriate to achieve strategic objectives.
- *External performance-based targets* are set to reflect the performance that is achieved by similar, or competitor, external operations.
- *Absolute performance targets* are based on the theoretical upper limit of performance

One of the problems in setting targets is that different targets can give very different messages regarding the improvement being achieved. So, for example, in Figure 12.4, one of an operation's performance measures is 'delivery' (in this case defined as the proportion of orders delivered on time). The performance for one month has been measured at 83 per cent, but any judgement regarding performance will be dependent on the performance targets. Using a *historical* target, when compared to last year's performance of 60 per cent, this months' performance of 83 per cent is good. But, if the operation's *strategy* calls for a 95 per cent delivery performance, the actual performance of 83 per cent looks decidedly poor. The company may also be concerned with how they perform against *competitors*' performance. If competitors are currently averaging delivery performances of around 80 per cent, the company's performance look rather good. Finally, the more ambitious managers within the company may wish to at least try and seek perfection. Why not, they argue, use an *absolute* performance

OPERATIONS PRINCIPLE
Performance measures only have meaning when compared against targets.

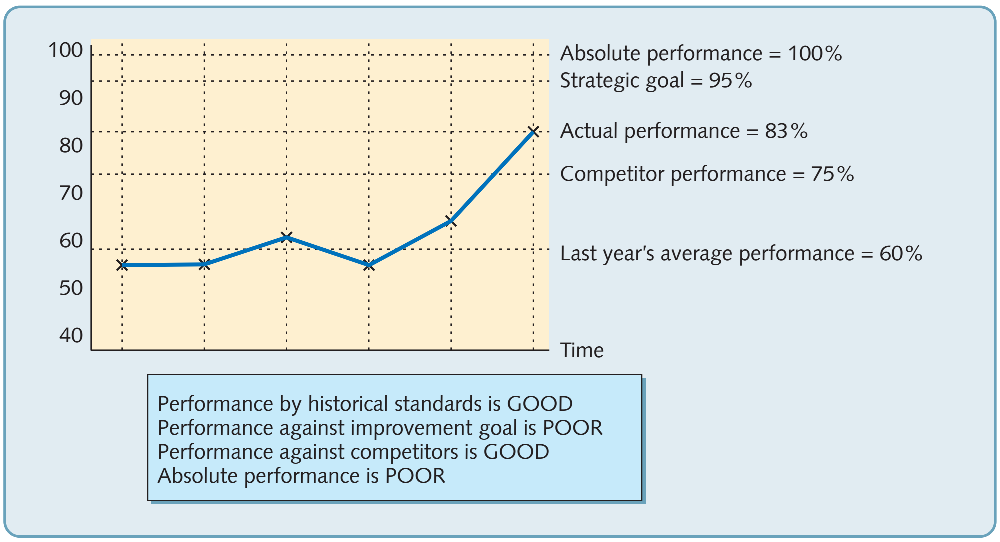

Figure 12.4 **Different standards of comparison give different messages**

standard of 100 per cent delivery on time? Against this standard the company's actual 83 per cent again looks disappointing.

Benchmarking

Benchmarking, is 'the process of learning from others' and involves comparing one's own performance, or methods against other comparable operations. It is a broader issue than setting performance targets, and includes investigating other organisations' operations practice in order to derive ideas that could contribute to performance improvement. Its rationale is based on the idea that (a) problems in managing processes are almost certainly shared by processes elsewhere, and (b) that there is probably another operation somewhere that has developed a better way of doing things. For example, a bank might learn some things from a supermarket about how it could cope with demand fluctuations during the day. Benchmarking is essentially about stimulating creativity in improvement practice.

OPERATIONS PRINCIPLE

Improvement is aided by contextualising processes and operations.

Types of benchmarking

There are many different types of benchmarking (which are not necessarily mutually exclusive), some of which are listed below:

- *Internal benchmarking* is a comparison between operations or parts of operations that are within the same total organisation. For example, a large motor vehicle manufacturer with several factories might choose to benchmark each factory against the others.
- *External benchmarking* is a comparison between an operation and other operations that are part of a different organisation.
- *Non-competitive benchmarking* is benchmarking against external organisations that do not compete directly in the same markets.
- *Competitive benchmarking* is a comparison directly between competitors in the same, or similar, markets.
- *Performance benchmarking* is a comparison between the levels of achieved performance in different operations. For example, an operation might compare its own performance in terms of some or all of our performance objectives – quality, speed, dependability, flexibility and cost – against other organisations' performance in the same dimensions.
- *Practice benchmarking* is a comparison between an organisation's operations practices, or way of doing things, and those adopted by another operation. For example, a large retail store might compare its systems and procedures for controlling stock levels with those used by another department store.

Benchmarking as an improvement tool

Although benchmarking has become popular, some businesses have failed to derive maximum benefit from it. Partly this may be because there are some misunderstandings about what benchmarking actually entails. First, it is not a 'one-off' project. It is best practised as a continuous process of comparison. Second, it does not provide 'solutions'. Rather, it provides ideas and information that can lead to solutions. Third, it does not involve simply copying or imitating other operations. It is a process of learning and adapting in a pragmatic manner. Fourth, it means devoting resources to the activity. Benchmarking cannot be done without some investment, but this does not necessarily mean allocating exclusive responsibility to a set of highly paid managers. In fact, there can be advantages in organising staff at all levels to investigate and collate information from benchmarking targets. There are also some basic rules about how benchmarking can be organised:

- A prerequisite for benchmarking success is to understand thoroughly your own processes. Without this it is difficult to compare your processes against those of other companies.

- Look at the information that is available in the public domain. Published accounts, journals, conferences and professional associations can all provide information which is useful for benchmarking purposes.
- Do not discard information because it seems irrelevant. Small pieces of information only make sense in the context of other pieces of information that may emerge subsequently.
- Be sensitive in asking for information from other companies. Don't ask any questions that we would not like to be asked ourselves.

EXAMPLE

Triumph motorcycles resurrected through benchmarking[6]

Triumph motorcycles once built the coolest bikes in the world. In the classic prisoner-of-war film, *The Great Escape,* Steve McQueen memorably jumped across the wire on a Triumph motorcycle. In the 1960s its larger motorcycles were selling well, both in the UK and America. But competition was catching up on the company, and just like the UK auto industry, Triumph declined from the 1970s onwards as better-designed and better-produced (mainly Japanese) products started to dominate the market. Within a few years the company had gone into receivership, and a property developer, John Bloor, bought the rights to the Triumph name relatively cheaply. He believed that there was a future for the company, yet he did not restart production immediately. Instead he spent years rethinking how the company's operations could be designed and run to compete in the modern motorcycle market. With his new team of managers, he went on an in-depth benchmarking study tour of Japan to analyse the production methods of those competitors who had driven the original Triumph into insolvency. '*We learned a lot*', said Nick Bloor, John's son, who now runs the company. It soon became clear to the management team that the original old factory in the West Midlands of the UK was not up to the task of producing world-class products. It was demolished and a new plant built in the UK that utilised the modern equipment and production methods learned on the Japanese visits. Now the company's plants in the UK and Thailand produce record numbers of bikes with styling that reflects the original bike's heritage, but with standards of engineering and reliability that match the operations that it learned from.

Deploying external ideas

Most of the literature that deals with improvement focuses on the generation, development and deployment of improvement ideas that originated within, rather from outside, the organisation. Yet to ignore the improvements that other companies are deploying is to ignore a potentially huge source of innovation. Whether they are competitors, suppliers, customers, or simply other firms with similar challenges, firms in the wider external business environment can provide solutions to internal problems. The discussion on benchmarking is clearly related to the idea of finding inspiration from outside the organisation. But some commentators on innovation go further and argue that (legally) 'copying' from outsiders can be an effective, if underused, approach to improvement. In his book, '*Copycats: How Smart Companies Use Imitation to Gain a Strategic Edge*'[7] Oded Shenkar claims that although, to argue 'imitation can be strategic seems almost blasphemous in the current scholarly climate', it can, 'be strategic and should be part of the strategic repertoire of any agile firm'. In fact, 'imitation can be a differentiating factor and has the potential to deliver unique value'. He cites Apple, making the point that the iPod was not the first digital-music player; nor was the iPhone the first smartphone, or the iPad the first tablet. To some extent Apple imitated ideas found in others' products, but solved the technical problems, established an appropriate supply chain operating model and made the products far more appealing. Similarly, Ray Kroc, who took McDonald's to worldwide success, copied White Castle, inventor of the fast-food burger joint. And Ireland's Ryanair imitated the business model originally developed in the US by Southwest Airlines.

Shenkar identifies three 'strategic types' of imitators:[8]

1. *The pioneer importer.* This is an imitator that is the pioneer in another place (another country, industry, or product market). This is what Ryanair did in Europe when it imported the Southwest model. Pioneer importer imitators may actually be able to move relatively slowly, especially if the original innovator, or other imitators, is unlikely to compete directly in the same market.
2. *The fast second.* This is a rapid mover arriving quickly after an innovator or pioneer, but before they have had an opportunity to establish an unassailable lead, and before other potentially rival imitators take a large share of the market. This strategy basically lets the pioneer take much of the risk of innovation in the hope that the follower can learn from the pioneer's experience.
3. *The come from behind.* This is a late entrant or adopter that has deliberately delayed adopting a new idea, maybe because of legal reasons, or because they want to be more certain that the idea will be acceptable. When they do adopt the idea, they may rely on differentiating themselves from the original pioneers. Samsung did this with their chip-making business, by using their manufacturing capability and knowledge to halve the time it takes to build a semiconductor plant. They then established a lead over competitors by exploiting their strength in key technical, production and quality skills.

EXAMPLE

Learning from Formula One[9]

As driving jobs go, there could be no bigger difference than between Formula One racing drivers weaving their way through some of the fastest competitors in the world and a supermarket truck driver quietly delivering beans, beer and bacon to distribution centres and stores. But they have more in common than one would suspect. Both Formula One and truck drivers want to save fuel, either to reduce pit-stops (Formula 1) or keep delivery costs down (heavy goods vehicles). And although grocery deliveries in the suburbs do not seem as thrilling as racing round the track at Monza, the computer-assisted simulation programs developed by the Williams Formula One team are being deployed to help Sainsbury's (a British supermarket group) drivers develop the driving skills that could potentially cut its fuel bill by up to 30 per cent. The simulator technology, which allows realistic advanced training to be conducted in a controlled environment, was developed originally for the advanced training of Formula One drivers and was developed and extended at the Williams Technology Centre in Qatar. It can now train drivers to a high level of professional driving skills and road safety applications.

Williams F1's chief executive, Alex Burns commented, '*Formula One is well recognised as an excellent technology incubator. It makes perfect sense to embrace some of the new and emerging technologies that the Williams Technology Centre in Qatar is developing from this incubator to help Sainsbury's mission to reduce its energy consumption and enhance the skills and safety of those supporting its crucial logistics operation.*' Sainsbury's energy-related improvement programmes tackle energy supply (for example, wind, solar and geothermal energy) as well as energy consumption (for example, switching to LED lighting, CO_2 refrigeration, etc.). Learning from Formula One will help Sainsbury's to improve further in the field of energy efficiency. Roger Burnley, Sainsbury's retail and logistics director, said, *'We are committed to reducing our environmental impact and, as a result, we are often at the very forefront of technological innovation. By partnering with Williams F1, we can take advantage of some of the world's most advanced automotive technology, making our operations even more efficient and taking us a step closer to meeting our CO_2 reduction targets.'*

Assess the gap between actual and target performance

A comparison of actual and target performance should guide the relative priorities for improvement. An important aspect of performance does not require immediate priority for improvement if current performance is already significantly better than target performance. Similarly, because some aspect of performance is relatively poor does not mean that it should be improved immediately if it exceeds its target performance. In fact, both the relative importance of the various performance measures, and their performance against target, need to be brought together in order to prioritise for improvement. One way of doing this is through the importance–performance matrix.

The importance–performance matrix

As its name implies, the importance–performance matrix positions each aspect of performance on a matrix according to its scores or ratings on how important each aspect of relative performance is, and what performance it is currently achieving. Figure 12.5 shows an importance–performance matrix divided into zones of improvement priority. The first zone boundary is the 'lower bound of acceptability' shown as line AB in Figure 12.5. This is the boundary between acceptable and unacceptable current performance. When some aspect of performance is rated as relatively unimportant, this boundary will be low. Most operations are prepared to tolerate lower performance for relatively unimportant performance factors. However, for performance factors that are rated more important, they will be markedly less sanguine at poor or mediocre levels of current performance. Below this minimum bound of acceptability (AB) there is clearly a need for

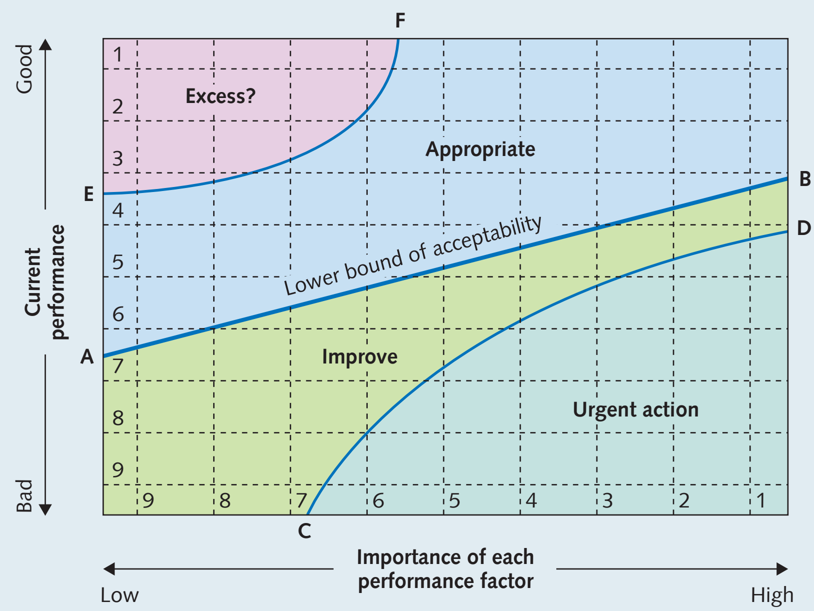

Figure 12.5 Priority zones in the importance–performance matrix

improvement; above this line, there is no immediate urgency for any improvement. However, not all factors of performance that fall below the minimum line will be seen as having the same degree of improvement priority. A boundary approximately represented by line CD represents a distinction between an urgent priority zone and a less urgent improvement zone. Similarly, above the line AB, not all competitive factors are regarded as having the same priority. The line EF can be seen as the approximate boundary between performance levels which are regarded as 'good' or 'appropriate' on the one hand and those regarded as 'too good' or 'excess' on the other. Segregating the matrix in this way results in four zones which imply very different priorities:

1. *The 'appropriate' zone.* Performance factors in this area lie above the lower bound of acceptability and so should be considered satisfactory.
2. *The 'improve' zone.* Lying below the lower bound of acceptability, any performance factors in this zone must be candidates for improvement.
3. *The 'urgent-action' zone.* These performance factors are important to customers but current performance is unacceptable. They must be considered as candidates for immediate improvement.
4. *The 'excess?' zone.* Performance factors in this area are 'high performing', but are not particularly important. The question must be asked, therefore, whether the resources devoted to achieving such a performance could be used better elsewhere.

EXAMPLE

EXL Laboratories

EXL Laboratories is a subsidiary of an electronics company. It carries out research and development as well as technical problem-solving work for a wide range of companies. It is particularly keen to improve the level of service that it gives to its customers. However, it needs to decide which aspect of its performance to improve first. It has devised a list of the most important aspects of its service:

- *The quality of its technical solutions* – the perceived appropriateness by customers.
- *The quality of its communications with customers* – the frequency and usefulness of information.
- *The quality of post-project documentation* – the usefulness of the documentation which goes with the final report.
- *Delivery speed* – the time between customer request and the delivery of the final report.
- *Delivery dependability* – the ability to deliver on the promised date.
- *Delivery flexibility* – the ability to deliver the report on a revised date.
- *Specification flexibility* – the ability to change the nature of the investigation.
- *Price* – the total charge to the customer.

EXL assigns a rating to each of these performance factors, both for their relative importance and their current performance, as shown in Figure 12.6. In this case, EXL have used a 1 to 9 scale, where 1 is 'very important', or 'good'. Any type of scale can be used.

EXL Laboratories plotted the relative importance and current performance ratings it had given to each of its performance factors on an importance–performance matrix. This is shown in Figure 12.7. It shows that the most important aspect of performance – the ability to deliver sound technical solutions to its customers – falls comfortably within the appropriate zone. Specification flexibility and delivery flexibility are also in the appropriate zone, although only just. Both delivery speed and delivery dependability seem to be in need of improvement as each is below the minimum level of acceptability for their respective importance positions. However, two competitive factors, communications and cost/price, are clearly in need of immediate improvement. These two factors should therefore be assigned the most urgent priority for improvement. The matrix also indicates that the company's documentation could almost be regarded as 'too good'.

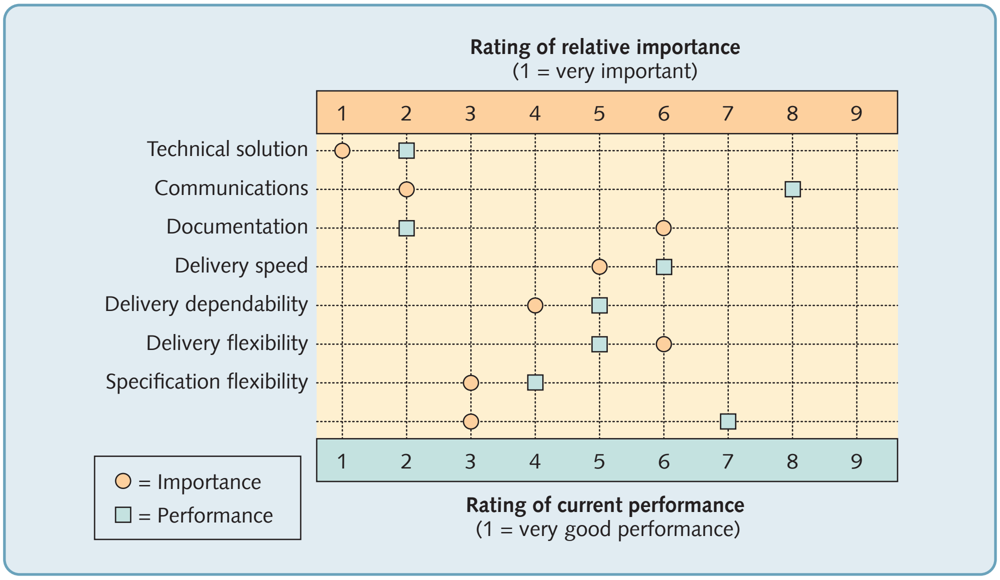

Figure 12.6 Rating 'importance' and 'current performance' for EXL Laboratories

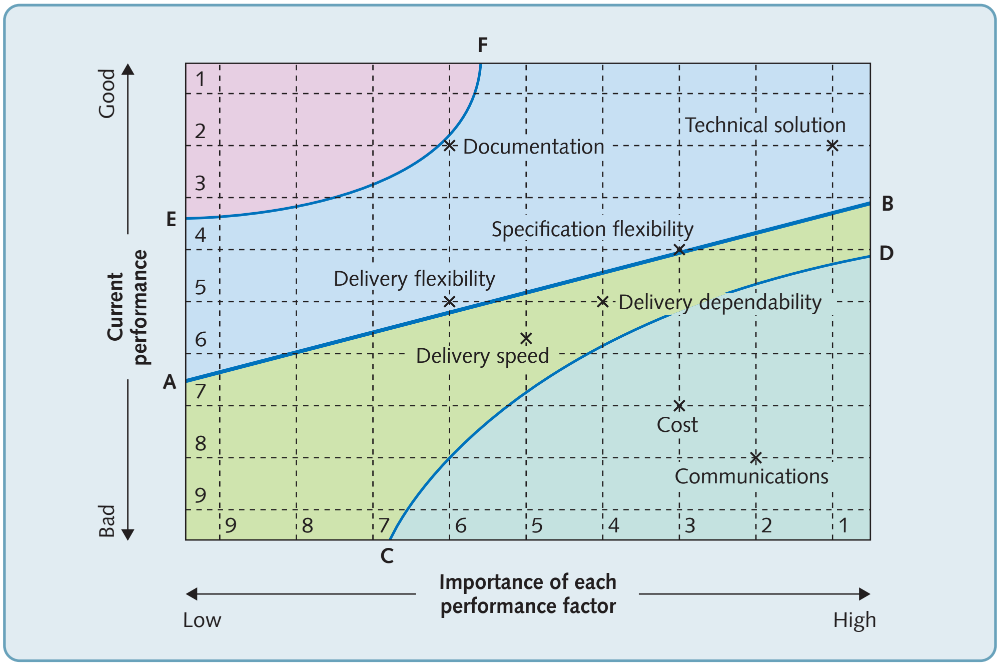

Figure 12.7 The importance–performance matrix for EXL Laboratories

The sandcone theory

As well as approaches that base improvement priority given on an operation's specific circumstances, some authorities believe that there is also a generic 'best' sequence of improvement. The best-known theory is called *the sandcone theory*,[10] so called because the sand is analogous to management effort and resources. Building a stable sandcone needs a stable foundation of quality, upon which one can build layers of dependability, speed, flexibility and cost (see Figure 12.8). Building up improvement is thus a cumulative process, not a sequential one. Moving on to the second priority for improvement does not mean dropping the first, and so on. According to the sandcone theory, the first priority should be *quality*, since this is a precondition to all lasting improvement. Only when the operation has reached a minimally acceptable

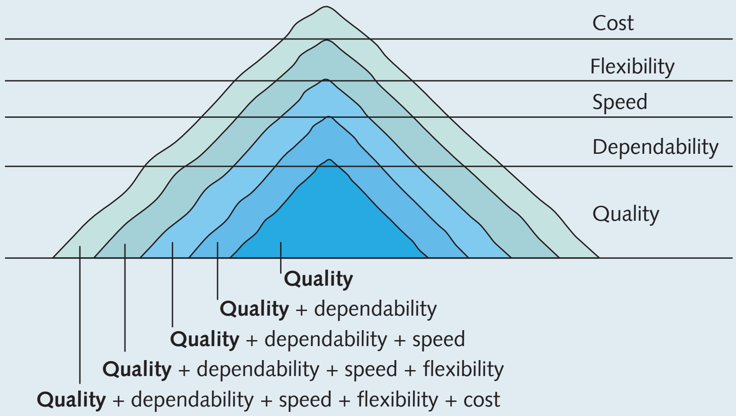

Figure 12.8 The sandcone model of improvement; cost reduction relies on a cumulative foundation of improvement in the other performance objectives

level in quality should it then tackle the next issue, that of internal *dependability*. Importantly though, moving on to include dependability in the improvement process will actually require further improvement in quality. Once a critical level of dependability is reached, enough to provide some stability to the operation, the next stage is to improve the *speed* of internal throughput. But again only while continuing to improve quality and dependability further. Soon it will become evident that the most effective way to improve speed is through improvements in response *flexibility*, that is, changing things within the operation faster. Again, including flexibility in the improvement process should not divert attention from continuing to work further on quality, dependability and speed. Only now, according to the sandcone theory, should *cost* be tackled head on.

DIAGNOSTIC QUESTION

What is the most appropriate improvement path?

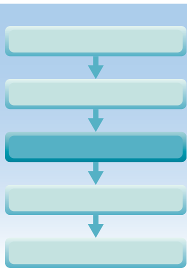

Once the priority of improvement has been determined, an operation must consider the approach or path it wishes to take to reaching its improvement goals. Two paths represent different, and to some extent opposing, philosophies; *breakthrough improvement* and *continuous improvement*, and although they represent different philosophies of improvement, they are not mutually exclusive. Few operations cannot benefit from improving their operations performance on a continuous basis, and few operations would reject investing in a major improvement breakthrough leap in performance if it represented good value. For most operations, both approaches are relevant to some extent, although possibly at different points in time. But to understand how and when each approach is appropriate one must understand their underlying philosophies.

OPERATIONS PRINCIPLE

Breakthrough and continuous improvement are not mutually exclusive.

Breakthrough improvement

Breakthrough (or 'innovation'-based) improvement assumes that the main vehicle of improvement is major and dramatic change in the way the operation works, for example, the total reorganisation of an operation's process structure, or the introduction of a fully integrated information system. The impact of these improvements represents a step-change in practice (and hopefully performance). Such improvements can be expensive, often disrupting the ongoing workings of the operation, and frequently involving changes in the product/service or process technology. Moreover, the actual level of improvement is not guaranteed. In fact, some improvement specialists argue that the greater the intended step-change in performance, the greater the risk that the actual increase in performance will disappoint.

The business process re-engineering approach

Typical of the radical breakthrough way of tackling improvement is the business process re-engineering (BPR) approach. It is a blend of a number of ideas such as fast throughput, waste elimination through process flow charting, customer-focused operations, and so on. But it was the potential of information technologies to enable the fundamental redesign of processes that acted as the catalyst in bringing these ideas together. BPR has been defined as:[11] *the fundamental rethinking and radical redesign of business processes to achieve dramatic improvements in critical, contemporary measures of performance, such as cost, quality, service and speed.*

Underlying the BPR approach is the belief that operations should be organised around the total process that adds value for customers, rather than the functions or activities that perform the various stages of the value-adding activity. (Figure 1.9 in Chapter 1 illustrates this idea.) The core of BPR is a redefinition of the processes within a total operation, to reflect the business processes that satisfy customer needs. Figure 12.9 illustrates this idea. The main principles of BPR have been summarised as follows:[12]

- Rethink business processes in a cross-functional manner that organises work around the natural flow of information (or materials or customers). This means organising around outcomes of a process rather than the tasks which go into it.
- Strive for dramatic improvements in the performance by radically rethinking and redesigning the process.
- Have those who use the output from a process perform the process. Check to see if all internal customers can be their own supplier, rather than depending on another function in the business to supply them (which takes longer and separates out the stages in the process).

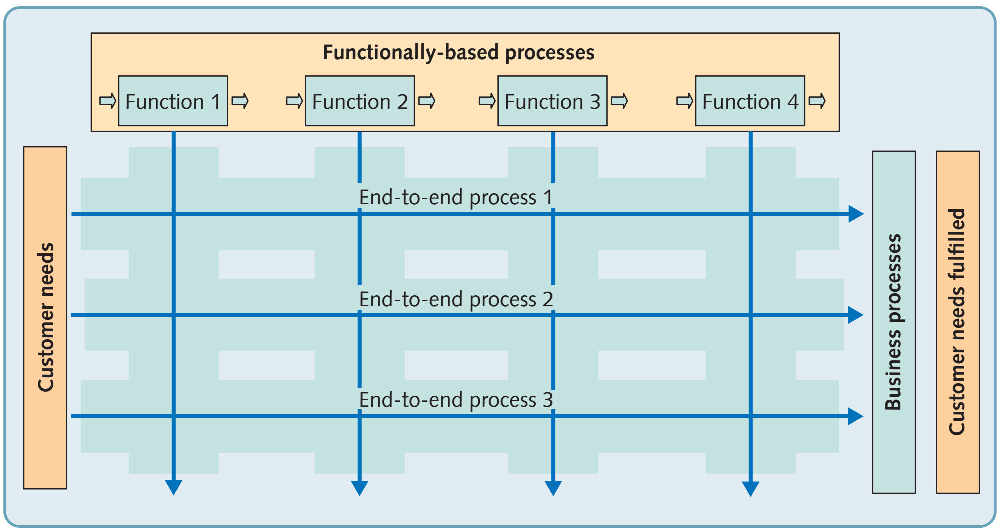

Figure 12.9 **BPR advocates reorganising (re-engineering) processes to reflect the natural processes that fulfil customer needs**

- Put decision points where the work is performed. Do not separate those who do the work from those who control and manage the work. Control and action are just one more type of supplier–customer relationship that can be merged.

Continuous improvement

Continuous improvement, as the name implies, adopts an approach to improving performance that assumes a never-ending series of small incremental improvement steps. For example, modifying the way a product is fixed to a machine to reduce changeover time, and simplifying the question sequence when taking a hotel reservation. While there is no guarantee that other steps will follow such small steps towards better performance, the whole philosophy of continuous improvement attempts to ensure that they will be. It is also known as *kaizen,* defined by Masaaki Imai (who has been one of the main proponents of continuous improvement) as follows: '*Kaizen means improvement. Moreover, it means improvement in personal life, home life, social life and work life. When applied to the work place, kaizen means continuing improvement involving everyone – managers and workers alike.*'

Continuous improvement is not concerned with promoting small improvements per se, but it does view small improvements as having one significant advantage over large ones – they can be followed relatively painlessly by others. It is not the *rate* of improvement which is important; it is the *momentum* of improvement. It does not matter if successive improvements are small; what does matter is that every month (or week, or quarter, or whatever period is appropriate) some kind of improvement has actually taken place. Continuous improvement does not always come naturally. There are specific abilities, behaviours and actions that need to be consciously developed if continuous improvement is to be sustained over the long term.

Improvement cycle models

An important element of continuous improvement is the idea that improvement can be represented by a never-ending process of repeatedly questioning and re-questioning the detailed working of a process. This is usually summarised by the idea of the *improvement cycle,* of which there are many, some proprietary models owned by consultancy companies. Two of the more generally used models are: the PDCA cycle (sometimes called the Deming cycle, named after the famous quality 'guru', W.E. Deming): and the DMAIC cycle (made popular by the Six-Sigma approach to improvement, see later).

OPERATIONS PRINCIPLE

Continuous improvement necessarily implies a never-ending cycle of analysis and action.

The PDCA (or PDSA) cycle

The PDCA cycle model is shown in Figure 12.10(a). It starts with the P (for plan) stage, which involves an examination of the current method or the problem area being studied. This involves collecting and analysing data to formulate a plan of action which is intended to improve performance. (Some of the techniques used to collect and analyse data are explained later.) The next step is the D (for do) stage. This is the implementation stage during which the plan is tried out in the operation. This stage may itself involve a mini-PDCA cycle as the problems of implementation are resolved. Next comes the C (for check) stage where the newly implemented solution is evaluated to see whether it has resulted in the expected improvement. Some versions of this idea use the term 'study' instead of 'check' and call the idea the 'PDSA' cycle, but the idea is basically the same. Finally, at least for this cycle, comes the A (for act) stage. During this stage the change is consolidated or standardised if it has been successful. Alternatively, if the change has not been successful, the lessons learned from the 'trial' are formalised before the cycle starts again. You may also find this cycle called the Deming cycle, Deming wheel or Shewhart cycle.

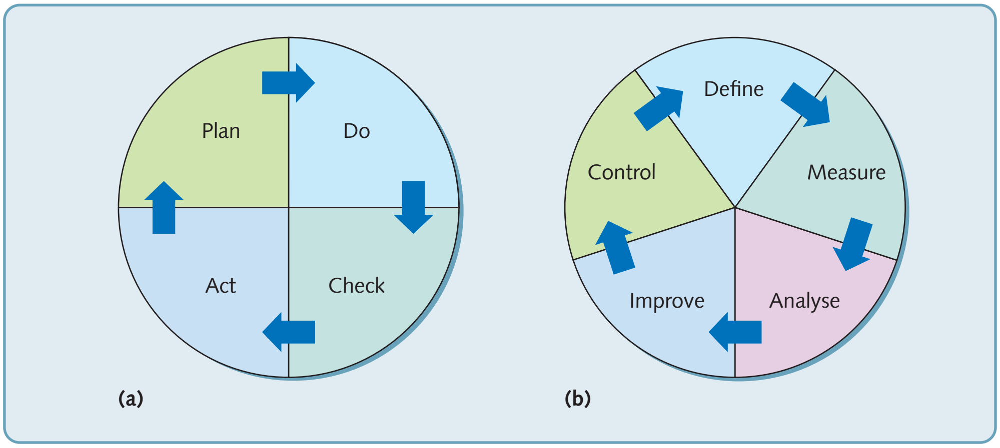

Figure 12.10 **(a) the plan-do-check-act, or 'Deming' improvement cycle and (b) the define-measure-analyse-improve-control, or DMAIC Six-Sigma improvement cycle**

The DMAIC cycle

In some ways this cycle is more intuitively obvious than the PDCA cycle in so much as it follows a more 'experimental' approach. The DMAIC cycle starts with defining the problem or problems, partly to understand the scope of what needs to be done and partly to define exactly the requirements of the process improvement. Often at this stage a formal goal or target for the improvement is set. After definition there is the measurement stage; this is important because the Six-Sigma approach emphasises the importance of working with hard evidence rather than opinion. It involves validating the problem (to make sure it is really worth solving), using data to refine the problem and measuring exactly what is happening. The analysis stage can be seen as an opportunity to develop hypotheses about what the root causes of the problem really are. Such hypotheses are validated (or not) by the analysis and the main root causes of the problem identified. Once the causes of the problem are identified, work can begin on improving the process. Ideas are developed to remove the root causes of problems, solutions are tested and those solutions that seem to work are implemented, formalised and the results are measured. The improved process needs then to be continually monitored and controlled to check that the improved level of performance is being sustained. The then cycle starts again, defining the problems that are preventing further improvement (see Figure 12.10(b)).

The last point in both cycles is the most important – *'the cycle starts again'*. It is only by accepting that in a continuous improvement philosophy these cycles quite literally never stop that improvement becomes part of every person's job.

The differences between breakthrough and continuous improvement

Breakthrough improvement places a high value on creative solutions, and encourages free thinking and individualism. It is a radical philosophy insomuch as it fosters an approach to improvement that does not accept many constraints on what is possible. 'Starting with a clean sheet of paper', 'going back to first principles' and 'completely rethinking the system' are all typical breakthrough improvement principles. Continuous improvement, on the other hand, is less ambitious, at least in the short-term. It stresses adaptability, teamwork and attention to detail. It is not radical; rather it builds upon the wealth of accumulated experience within the operation itself, often relying primarily on the people who operate the system to improve it. One analogy used to explain this difference is the sprint versus the marathon. Breakthrough improvement is a series of explosive and impressive sprints. Continuous improvement, like marathon running,

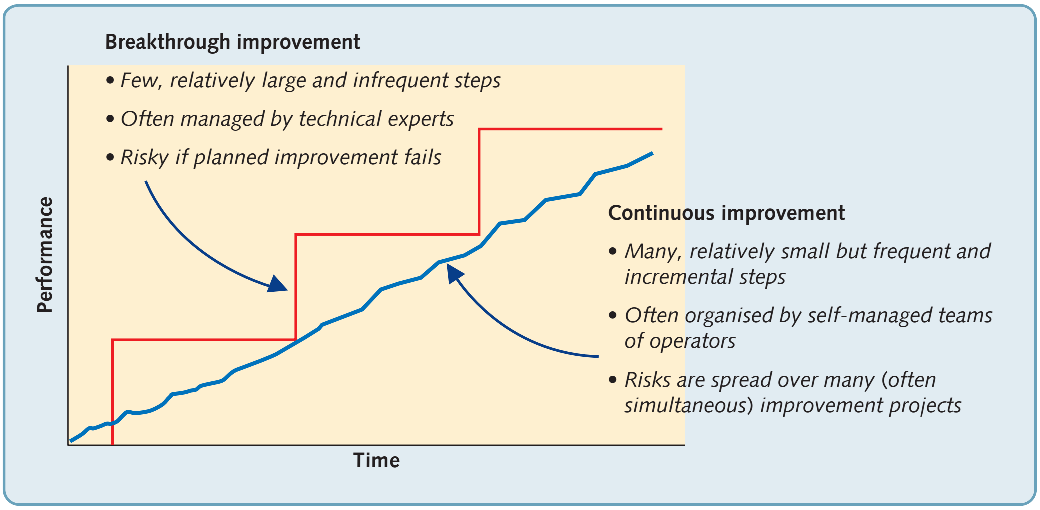

Figure 12.11 The differences between the breakthrough and continuous approaches to improvement

does not require the expertise and prowess required for sprinting; but it does require that the runner (or operations manager) keeps on going. Yet notwithstanding these differences, it is possible to use both approaches. Large and dramatic improvements can be implemented as and when they seem to promise significant improvement steps, but between such occasions, the operation can continue making its quiet and less spectacular kaizen improvements. Figure 12.11 shows some of the differences between the two approaches.

OPERATIONS PRINCIPLE
Breakthrough improvement necessarily implies radical and/or extensive change.

Exploitation or exploration

A closely related distinction to that between continuous and breakthrough improvement, is the one that management theorists draw between what they call 'exploitation' versus 'exploration'. Exploitation is the activity of enhancing processes (and products) that already exist within a firm. The focus of exploitation is on creating efficiencies rather than radically changing resources or processes. Its emphasis is on tight control of the improvement process, standardising processes, clear organisational structures and organisational stability. The benefits from exploitation tend to be relatively immediate, incremental and predictable. They also are likely to be better understood by the firm and fit into its existing strategic framework. Exploration, by contrast, is concerned with the exploration of new possibilities. It is associated with searching for and recognising new mindsets and ways of doing things. It involves experimentation, taking risks, simulation of possible consequences, flexibility and innovation. The benefits from exploration are principally long-term but can be relatively difficult to predict. Moreover, any benefits or discoveries that might come may be so different from what the firm is familiar with that it may not find it easy to take advantage of them.

Organisational 'ambidexterity'

It is clear that the organisational skills and capabilities to be successful at exploitation are likely to be very different from those that are needed for the radical exploration of new ideas. Indeed, the two views of improvement may actively conflict. A focus on thoroughly exploring for totally novel choices may consume managerial time, effort and the financial resources that would otherwise be used for refining existing ways of doing things, reducing the effectiveness of improving existing processes. Conversely, if existing processes are improved over time, there may be less motivation to experiment with new ideas. So, although both exploitation and exploration can be beneficial, they may compete both for resources and for management attention. This is where the concept of 'organisational ambidexterity' becomes important. Organisational ambidexterity means the ability of a firm to both exploit and explore as they seek to improve; to be able to compete in mature markets where efficiency is important, by improving existing resources and processes, while also competing in new technologies and/or markets where novelty, innovation and experimentation are required.

The Six Sigma approach to organising improvement

One approach to improvement that combines breakthrough and continuous philosophies is *Six Sigma.* Although technically the 'Six Sigma' name derives from statistical process control (SPC), and more specifically the concept of process capability, it has now come to mean a much broader approach to improvement. The following definition gives a sense of its modern usage. *'Six Sigma is a comprehensive and flexible system for achieving, sustaining and maximising business success. Six Sigma is uniquely driven by close understanding of customer needs, disciplined used of facts, data, and statistical analysis, and diligent attention to managing, improving, and reinventing business processes.'*[13]

The Six Sigma concept, therefore, includes many of the issues covered in this and other chapters of this book; for example, process design and redesign, balanced scorecard measures, continuous improvement, statistical process control, ongoing process planning and control, and so on. However, at the heart of Six Sigma lies an understanding of the negative effects of variation in all types of business process. This aversion to variation was first popularised by Motorola, the electronics company, who set its objective as 'total customer satisfaction' in the 1980s, then decided that true customer satisfaction would only be achieved when its products were delivered when promised, with no defects, with no early-life failures and no excessive failure in service. To achieve this, they initially focused on removing manufacturing defects, but it soon realised that many problems were caused by latent defects, hidden within the design of its products. The only way to eliminate these defects was to make sure that design specifications were tight (i.e. narrow tolerances) and its processes very capable.

Motorola's Six-Sigma quality concept was so named because it required that the natural variation of processes ($\pm$ 3 standard deviations) should be half their specification range. In other words, the specification range of any part of a product or service should be $\pm$ 6 the standard deviation of the process. The Greek letter sigma (σ) is often used to indicate the standard deviation of a process, hence the Six-Sigma label. The Six Sigma approach also uses the measure of 'defects per million *opportunities*' (DPMO). This is the number of defects that the process will produce if there were one million opportunities to do so. So, difficult processes with many opportunities for defects can be compared with simple processes with few opportunities for defects.

The Six Sigma approach also holds that improvement initiatives can only be successful if significant resources and training are devoted to their management. It recommends a specially trained cadre of practitioners, many of whom should be dedicated full time to improving processes as internal consultants. The terms that have become associated with this group of experts (and denote their level of expertise) are Master Black Belt, Black Belt and Green Belt:

- *Master Black Belts* are experts in the use of Six Sigma tools and techniques, as well as how such techniques can be used and implemented. They are seen as teachers who can not only guide improvement projects, but also coach and mentor Black Belts and Green Belts. Given their responsibilities, it is expected that Master Black Belts are employed full time on their improvement activities.
- *Black Belts* take a direct hand in organising improvement teams, and will usually have undertaken a minimum of 20 to 25 days training and carried out at least one major improvement

project. Black Belts are expected to develop their quantitative analytical skills and also act as coaches for Green Belt. Like Master Black Belts, they are dedicated full time to improvement, and although opinions vary, some organisations recommend one Black Belt for every 100 employees.

- *Green Belts* work within improvement teams, possibly as team leaders. They have less training than Black Belts – typically around 10 to 15 days. Green Belts are not full-time positions. They have normal day-to-day process responsibilities but are expected to spend at least 20 per cent of their time on improvement projects.

Devoting such large amounts of training and time to improvement is a significant investment, especially for small companies. Nevertheless, Six Sigma proponents argue that the improvement activity is generally neglected in most operations and if it is to be taken seriously, it deserves the significant investment implied by the Six Sigma approach. Furthermore, they argue, if operated well, Six Sigma improvement projects run by experienced practitioners can save far more than their cost.

DIAGNOSTIC QUESTION

What techniques should be used to facilitate improvement?

OPERATIONS PRINCIPLE

Improvement is facilitated by relatively simple analytical techniques.

All the techniques described in this book and its supplements can be regarded as 'improvement' techniques. However, some techniques are particularly useful for improving operations and processes generally. Here we select some techniques that either have not been described elsewhere or need to be reintroduced in their role of helping operations improvement.

Scatter diagrams

Scatter diagrams provide a quick and simple method of identifying whether there is evidence of a connection between two sets of data: for example, the time at which you set off for work every morning and how long the journey to work takes. Plotting each journey on a graph, which has departure time on one axis and journey time on the other, could give an indication of whether departure time and journey time are related and, if so, how. Scatter diagrams can be treated in a far more sophisticated manner by quantifying how strong the relationship between the sets of data is. But, however sophisticated the approach is, this type of graph only identifies the existence of a relationship, not necessarily the existence of a cause–effect relationship. If the scatter diagram shows a very strong connection between the sets of data, it is important evidence of a cause–effect relationship, but not proof positive. It could be coincidence!

EXAMPLE

Kaston Pyral Services Ltd (1)

Kaston Pyral Services Ltd (KPS) installs and maintains environmental control, heating and air conditioning systems. It has set up an improvement team to suggest ways in which it might improve its levels of customer service. The improvement team had completed its first customer satisfaction survey. The survey asked customers to score the service they received from KPS in several ways. For example, it asked customers to score services on a scale of 1 to 10 on promptness, friendliness, level of advice, and so on. Scores were then summed to give a 'total

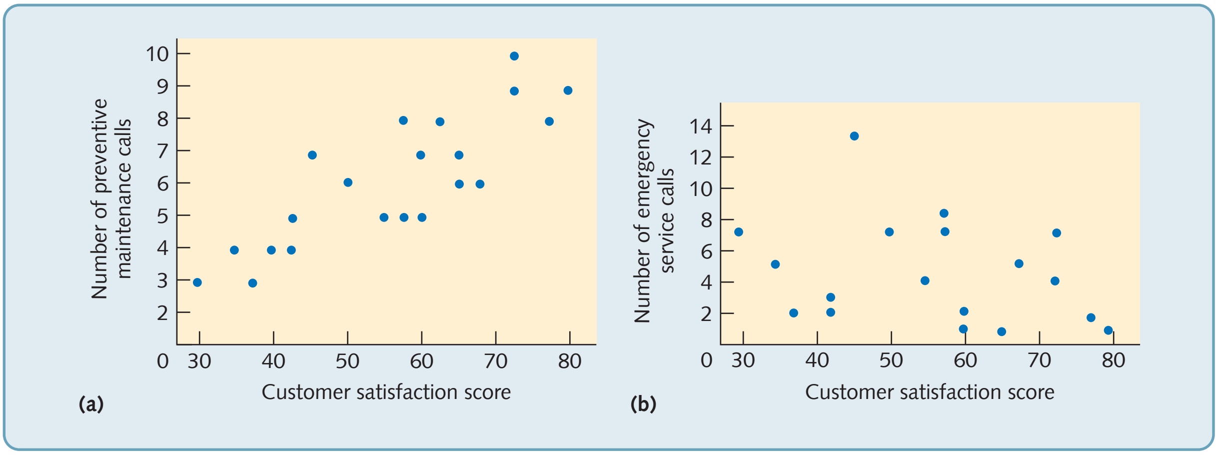

Figure 12.12 Scatter diagrams for customer satisfaction versus (a) number of preventive maintenance calls and (b) number of emergency service calls

satisfaction score' for each customer – the higher the score, the greater the satisfaction. The spread of satisfaction scores puzzled the team and they considered what factors might be causing such differences in the way their customers viewed them. Two factors were put forward to explain the differences:

1. the number of times in the past year the customer had received a preventive maintenance visit
2. the number of times the customer had called for emergency service.

All these data were collected and plotted on scatter diagrams as shown in Figure 12.12. Figure 12.12(a) shows that there seems to be a clear relationship between a customer's satisfaction score and the number of times the customer was visited for regular servicing. The scatter diagram in Figure 12.12(b) is less clear. Although all customers who had very high satisfaction scores had made very few emergency calls, so too had some customers with low satisfaction scores. As a result of this analysis, the team decided to survey customers' views on its emergency service.

Cause–effect diagrams

Cause–effect diagrams are a particularly effective method of helping to search for the root causes of problems. They do this by asking what, when, where, how and why questions, but also add some possible 'answers' in an explicit way. They can also be used to identify areas where further data are needed. Cause–effect diagrams (which are also known as Ishikawa diagrams) have become extensively used in improvement programmes. This is because they provide a way of structuring group brainstorming sessions. Often the structure involves identifying possible causes under the (rather old-fashioned) headings of: machinery, manpower, materials, methods and money. Yet in practice, any categorisation that comprehensively covers all relevant possible causes could be used.

EXAMPLE

Kaston Pyral Services Ltd (2)

The improvement team at KPS was working on a particular area that was proving to be a problem. Whenever service engineers were called out to perform emergency servicing for a customer, they took with them the spares and equipment that they thought would be necessary to repair the system. Although engineers could never be sure exactly what materials and equipment they would need for a job, they could guess what was likely to be needed and take

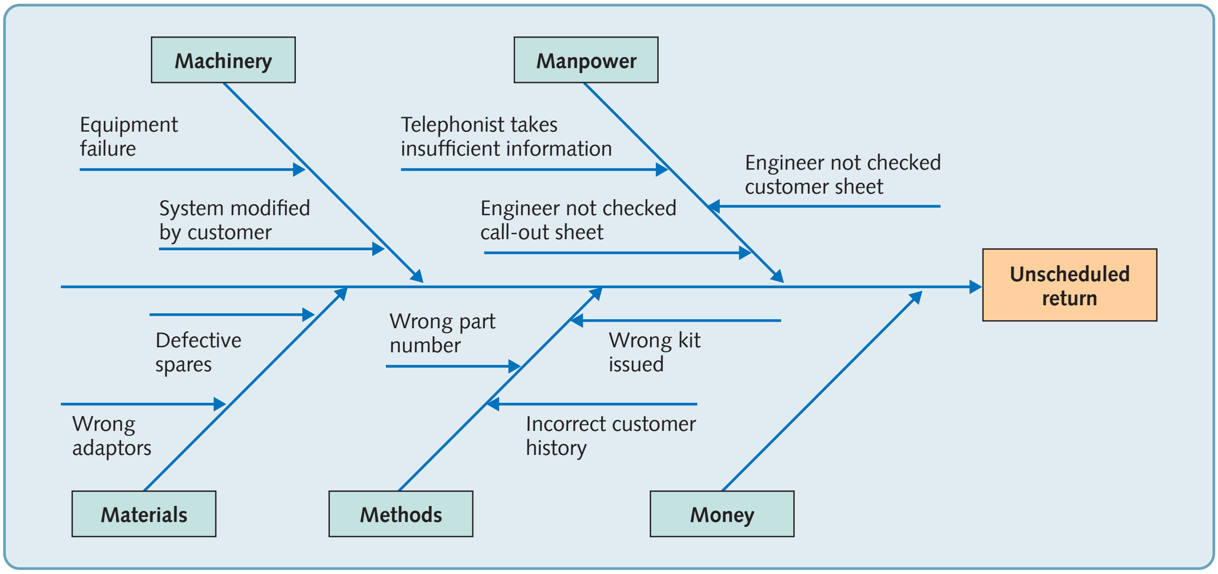

Figure 12.13 **Cause–effect diagram of unscheduled returns at KPS**

a range of spares and equipment which would cover most eventualities. Too often, however, the engineers would find that they needed a spare that they had not brought with them. The cause–effect diagram for this particular problem, as drawn by the team, is shown in Figure 12.13.

Pareto diagrams

In any improvement process, it is worthwhile distinguishing between what is important and what is less so. The purpose of the Pareto diagram (that was first introduced in Chapter 9) is to distinguish between the 'vital few' issues and the 'trivial many'. It is a relatively straightforward technique that involves arranging items of information on the types of problem or causes of problem into their order of importance (usually measured by 'frequency of occurrence'). This can be used to highlight areas where further decision-making will be useful. Pareto analysis is based on the phenomenon of relatively few causes explaining the majority of effects. For example, most revenue for any company is likely to come from relatively few of the company's customers. Similarly, relatively few of a doctor's patients will probably occupy most of his or her time.

EXAMPLE

Kaston Pyral Services Ltd (3)

The KPS improvement team that was investigating unscheduled returns from emergency servicing (the issue described in the cause–effect diagram in Figure 12.13) examined all occasions over the previous 12 months on which an unscheduled return had been made. They categorised the reasons for unscheduled returns as follows:

1. The wrong part had been taken to a job because, although the information which the engineer received was sound, he or she had incorrectly predicted the nature of the fault.
2. The wrong part had been taken to the job because there was insufficient information given when the call was taken.
3. The wrong part had been taken to the job because the system had been modified in some way not recorded on KPS's records.
4. The wrong part had been taken to the job because the part had been incorrectly issued to the engineer by stores.

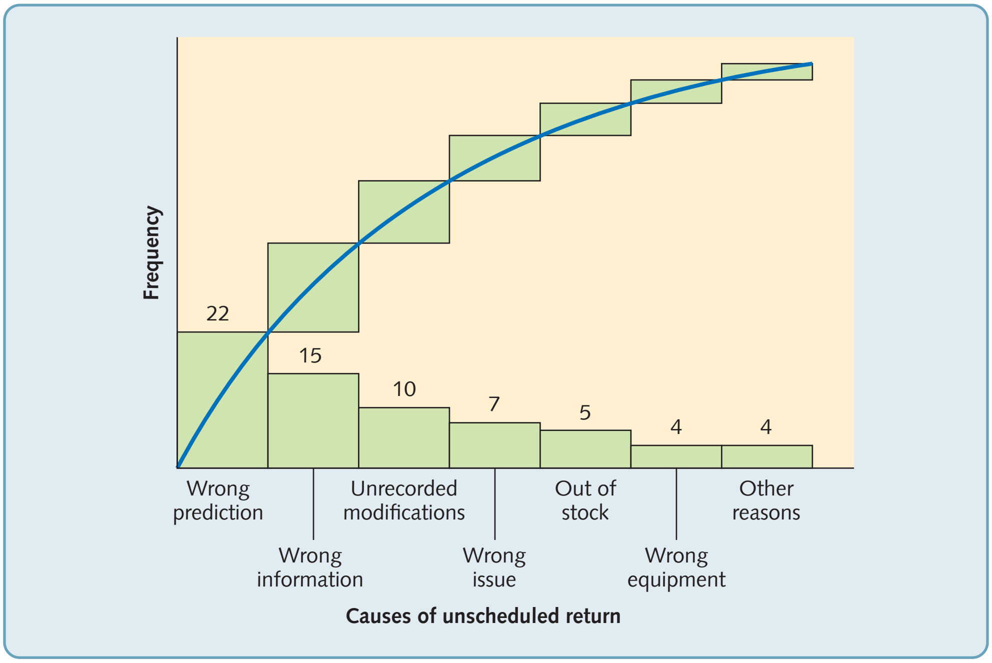

Figure 12.14 **Pareto diagram for causes of unscheduled returns**

5. No part had been taken because the relevant part was out of stock.
6. The wrong equipment had been taken for whatever reason.
7. Any other reason.

The relative frequency of occurrence of these causes is shown in Figure 12.14. About a third of all unscheduled returns were due to the first category, and more than half the returns were accounted for by the first and second categories together. It was decided that the problem could best be tackled by concentrating on how to get more information to the engineers that would enable them to predict the causes of failure accurately.

Why–why analysis

Why–why analysis starts by stating the problem and asking *why* that problem has occurred. Once the major reasons for the problem occurring have been identified, each one is taken in turn and again the question is asked *why* those reasons have occurred, and so on. This procedure is continued until either a cause seems sufficiently self-contained to be addressed by itself or no more answers to the question 'Why?' can be generated.

EXAMPLE

Kaston Pyral Services Ltd (4)

The major cause of unscheduled returns at KPS was the incorrect prediction of reasons for the customer's system failure. This is stated as the 'problem' in the why–why analysis in Figure 12.15. The question is then asked, why was the failure wrongly predicted? Three answers are proposed: first, the engineers were not trained correctly; second, they had insufficient knowledge of the particular product installed in the customer's location; and third, they had insufficient knowledge of the customer's particular system with its modifications. Each of these three reasons is taken in turn, and the questions are asked, why is there a lack of training, why is there a lack of product knowledge and why is there a lack of customer knowledge? And so on.

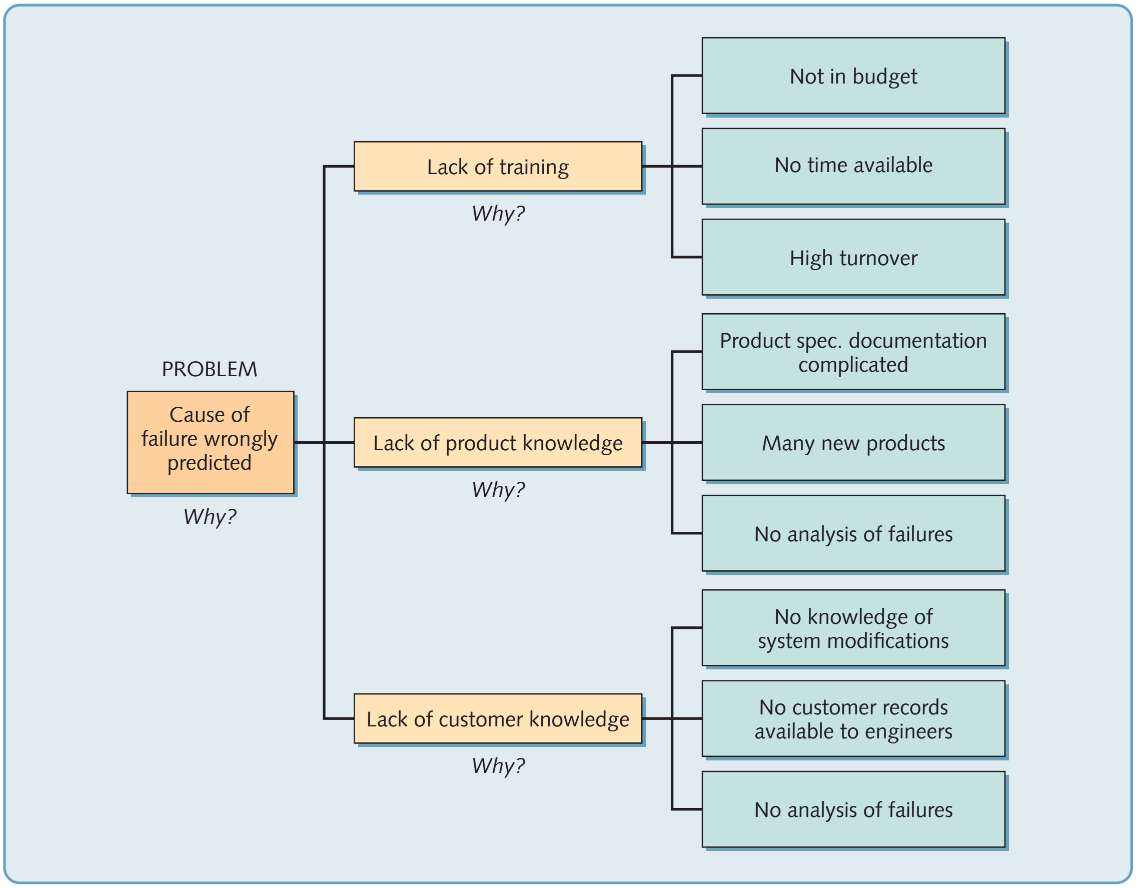

Figure 12.15 Why–why analysis for 'failure wrongly predicted'

DIAGNOSTIC QUESTION

How can improvement be made to stick?

Not all of the improvement initiatives (often launched with high expectations), will go on to fulfil their potential. Even those improvement initiatives that are successfully implemented may lose impetus over time. Sometimes this is because of managers' view of the nature of improvement, and at other times it is because managers fail to manage the improvement process adequately. Nor does a successful improvement culture necessarily rely on the kind of techniques described earlier in this chapter (although they clearly have a role). Sometimes the simplest approaches can be remarkably effective, as in the idea of 'checklists' described next.

EXAMPLE

The Checklist Manifesto[14]

Improvement methodologies are often associated with repetitive operations. Performing the same task repeatedly means that there are plenty of opportunities to 'get it right'. The whole idea behind continuous improvement derives for this simple idea. By contrast, operations that have to perform more difficult activities, especially those that call for expert judgement and diagnostic ability must call for equally complex improvement approaches. No? Well 'no', according to Atul Gawande, a physician at the prestigious Johns Hopkins Hospital. Mr Gawande thinks that the very opposite is true. Although medicine is advancing at an astounding rate and medical journals produce learned papers adding the results of advanced research to an ever-expanding pool of knowledge. Surgeons carry out over 200 major operations a year, unfortunately not all of them successful, but the medical profession overall does not always have a reliable method for learning from its mistakes. Atul Gawande's idea is that his, and similar 'knowledge-based' professions are in danger of sinking under the weight of facts. Scientists are accumulating more and more information and professions are fragmenting into ever-narrower specialisms. Mr Gawande tells the story of Peter Pronovost, a specialist in critical care at Johns Hopkins Hospital who tried to reduce the number of patients that were becoming infected on account of the use of intravenous central lines. There are five steps that medical teams can take to reduce the chances of contracting such infections. Initially, Pronovost simply asked nurses to observe whether doctors took the five steps. What they found was that, at least a third of the time, they missed one or more of the steps. So nurses were authorised to stop doctors who had missed out any of the steps, and, as a matter of course, ask whether existing intravenous central lines should be reviewed. As a result of applying these simple checklist style rules, the ten-day line-infection rates went down from 11 per cent to zero. In one hospital, it was calculated that, over a year, this simple method had prevented 43 infections, 8 deaths and saved about $2 million. Using the same checklist approach, the hospital identified and applied the method to other activities. For example, a check in which nurses asked patients about their pain levels led to untreated pain reducing from 41 per cent to 3 per cent. Similarly, the simple checklists method helped the average length of patient stay in intensive care to fall by half. When Pronovost's approach was adopted by other hospitals, within 18 months, 1,500 lives and $175 million had been saved.

Mr Gawande describes checklists used in this way as a 'cognitive net' – a mechanism that can help prevent experienced people from making errors due to flawed memory and attention, and ensure that teams work together. Simple checklists are common in other professions. Civil engineers use them to make certain that complicated structures are assembled on schedule. Chefs use them to make sure that food is prepared exactly to the customers' taste. Airlines use them to make sure that pilots take-off safely and also to learn from, now relatively rare, crashes. Indeed, Mr Gawande is happy to acknowledge that checklists are not a new idea. He tells the story of the prototype of the Boeing B17 Flying Fortress that crashed after take-off on its trial flight in 1935. Most experts said that the bomber was 'too complex to fly'. Facing bankruptcy, Boeing, investigated and discovered that, confronted with four engines rather than two, the pilot forgot to release a vital locking mechanism. But Boeing created a pilot's checklist, in which the fundamental actions for the stages of flying were made a mandated part of the pilot's job. In the following years, B17s flew almost two million miles without a single accident. Even for pilots, many of whom are rugged individualists, says Mr Gawande, it is usually the application of routine procedures that saves planes when things go wrong, rather than 'hero-pilotry', so fêted by the media. It is discipline rather than brilliance that preserves life. In fact, it is discipline that leaves room for brilliance to flourish.

Avoid becoming a victim of improvement 'fashion'

Improvement has, to some extent, become a fashion industry with new ideas and concepts continually being introduced as offering a novel way to improve business performance. There is nothing intrinsically wrong with this. Fashion stimulates and refreshes through introducing novel ideas. Without it, things would stagnate. The problem lies not with new improvement

OPERATIONS PRINCIPLE
The popularity of an improvement approach is not necessarily an indicator of its effectiveness.

ideas, but rather with some managers becoming a victim of the process, where some new idea will entirely displace whatever went before. Most new ideas have something to say, but jumping from one fad to another, will not only generate a backlash against any new idea, but also destroy the ability to accumulate the experience that comes from experimenting with each one.

Avoiding becoming an improvement fashion victim is not easy. It requires that those directing the improvement process take responsibility for a number of issues:

- They must take responsibility for improvement as an ongoing activity, rather than becoming champions for only one specific improvement initiative.
- They must take responsibility for understanding the underlying ideas behind each new concept. Improvement is not 'following a recipe' or 'painting by numbers'. Unless one understands *why* improvement ideas are supposed to work, it is difficult to understand *how* they can be made to work properly.
- They must take responsibility for understanding the antecedents to a 'new' improvement idea, because it helps to understand it better and to judge how appropriate it may be for one's own operation.
- They must be prepared to adapt new ideas so that they make sense within the context of their own operation. 'One size' rarely fits all.
- They must take responsibility for the (often significant) education and learning effort that will be needed if new ideas are to be intelligently exploited.
- Above all, they must avoid the over-exaggeration and hype that many new ideas attract. Although it is sometimes tempting to exploit the motivational 'pull' of new ideas through slogans, posters and exhortations, carefully thought out plans will always be superior in the long run, and will help avoid the inevitable backlash that follows 'over-selling' a single approach.

Managing the improvement process

There is no absolute prescription for the way improvement should be managed. Any improvement process should reflect the uniqueness of each operation's characteristics. What appears to be almost a guarantee of difficulty in managing improvement processes, are attempts to squeeze improvement into a standard mould. Nevertheless, there are some aspects of any improvement process that appear to influence its eventual success, and should at least be debated.

OPERATIONS PRINCIPLE
There is no one universal approach to improvement.

Should an improvement *strategy* be defined?

Without thinking through the overall purpose and long-term goals of the improvement process it is difficult for any operation to know where it is going. Specifically, an improvement strategy should have something to say about:

- the competitive priorities of the organisation, and how the improvement process is expected to contribute to achieving increased strategic impact
- the roles and responsibilities of the various parts of the organisation in the improvement process
- the resources which will be available for the improvement process
- the general approach to, and philosophy of, improvement in the organisation.

Yet, too rigid a strategy can become inappropriate if the business's competitive circumstances change, or as the operation learns through experience. But, the careful modification of improvement strategy in the light of experience is not the same as making dramatic changes in improvement strategy as new improvement fashions appear.

What degree of top-management support is required?

For most authorities, the answer is unambiguous – a significant amount. Without top- management support, improvement cannot succeed. It is the most crucial factor in almost all the studies of improvement process implementation. It also goes far beyond merely allocating senior resources to the process. 'Top-management support' usually means that senior personnel must:

- understand and believe in the link between improvement and the business's overall strategic impact
- understand the practicalities of the improvement process and be able to communicate its principles and techniques to the rest of the organisation
- be able to participate in the total problem-solving process to improve performance
- formulate and maintain a clear idea of the operation's improvement philosophy.

Should the improvement process be formally supervised?

Some improvement processes fail because they develop an unwieldy 'bureaucracy' to run them. But any process needs to be managed, so all improvement processes will need some kind of group to design, plan and control its efforts. However, a worthwhile goal for many improvement processes is to make it 'self-governing' over time. In fact, there are significant advantages in terms of people's commitment in giving them responsibility for managing the improvement process. However, even when improvement is driven primarily by self-managing improvement groups, there is a need for some sort of 'repository of knowledge' to ensure that the learning and experience accumulated from the improvement process is not lost.

To what extent should improvement be group-based?

No one can really know a process quite like the people who operate it. They have access to the informal as well as the formal information networks that contain the way processes really work. But, working alone, individuals cannot pool their experience or learn from one another. So improvement processes are almost always based on teams. The issue is how these teams should be formulated, which will depend on the circumstances of the operation, its context and its objectives. For example, *quality circles,* much used in Japan, encountered mixed success in the West. A very different type of team is the '*task force*', or what some US companies call a 'tiger team'. Compared with quality circles, this type of group is far more management directed and focused. Most improvement teams are between these two extremes.

How should success be recognised?

If improvement is so important, it should be recognised, with success, effort and initiative being formally rewarded. The paradox is that if improvement is to become part of everyday operational life, then why should improvement effort be especially rewarded? One compromise is to devise a recognition-and-rewards system that responds to improvement initiatives early in the improvement process, but then merges into the operation's normal reward procedures. In this way, people are rewarded not just for the efficient and effective running of their processes on an ongoing basis, but also for improving their processes. Then improvement will become an everyday responsibility of all people in the operation.

How much training is required?

Training has two purposes in the development of improvement processes. The first is to provide the necessary skills that will allow staff to solve process problems and implement improvements. The second is to provide an understanding of the appropriate interpersonal, group and organisational skills that are needed to 'lubricate' the improvement process. This second objective is more difficult than the first. Training and improvement techniques may take up significant time and effort, but none of this knowledge will be of much use if the organisational context for improvement mitigates against the techniques being used effectively. Although the nature of appropriate organisational development is beyond the scope of this book, it is worth noting both technique-based skills and organisational skills are enhanced if staff have a basic understanding of the core ideas and principles of operations and process management.

Critical commentary

Many of the issues covered in this chapter are controversial, for different reasons. Some criticism concerns the effectiveness of improvement methods. For example, it can be argued that there is a fundamental flaw in the concept of benchmarking. Operations that rely on others to stimulate their creativity, especially those that are in search of 'best practice', are always limiting themselves to currently accepted methods of operating or currently accepted limits to performance. 'Best practice' is not 'best' in the sense that it cannot be bettered, it is only 'best' in the sense that it is the best one can currently find. And accepting what is currently defined as 'best' may prevent operations from ever making the radical breakthrough or improvement that takes the concept of 'best' to a new and fundamentally improved level. Furthermore, because one operation has a set of successful practices in the way it manages its process does not mean that adopting those same practices in another context will prove equally successful. It is possible that subtle differences in the resources within a process (such as staff skills or technical capabilities) or the strategic context of an operation (for example, the relative priorities of performance objectives) will be sufficiently different to make the adoption of seemingly successful practices inappropriate.

- Other approaches are seen by some as too radical and too insensitive. For example, business process re-engineering has aroused considerable controversy. Most of its critics are academics, but some practical objections to BPR have also been raised, such as the fear that BPR looks only at work activities rather than at the people who perform the work. Because of this, people become 'cogs in a machine'. Also some see BPR as being too imprecise because its proponents cannot agree on whether it has to be radical or whether it can be implemented gradually, or exactly what a process is, or whether it has to be top-down or bottom-up, or on whether it has be supported by information technology or not. Perhaps most seriously, BPR is viewed as merely an excuse for getting rid of staff. Companies that wish to 'downsize' (that is, reduce numbers of staff within an operation) are using BPR as an excuse. This puts the short-term interests of the shareholders of the company above either their longer-term interests, or the interests of the company's employees. Moreover, a combination of radical redesign together with downsizing can mean that the essential core of experience is lost from the operation. This leaves it vulnerable to any marked turbulence, since it no longer has the knowledge and experience of how to cope with unexpected changes.

- Even the more gentle approach of continuous improvement is not universally welcomed. Notwithstanding its implications of empowerment and liberal attitude towards shop-floor staff, it is regarded by some worker representatives as merely a further example of management exploiting workers. Its critics have defined relatively established ideas such as TQM as 'management by stress'. Or, even more radically, 'TQM is like putting a vacuum cleaner next to a worker's brain and sucking out ideas. They don't want to rent your knowledge anymore, they want to own it – in the end that makes you totally replaceable.'

SUMMARY CHECKLIST

- ☐ Is the importance of performance improvement fully recognised within the operation?
- ☐ Do all operations and process managers see performance improvement as an integral part of their job?
- ☐ Is the gap between current and desired performance clearly articulated in all areas?
- ☐ Is the current performance measurement system seen as forming a basis for improvement?
- ☐ Does performance measurement focus on factors that reflect the operation's strategic objectives?
- ☐ Do performance measures allow likely problem areas to be diagnosed?
- ☐ Is some kind of balanced scorecard approach used that includes financial, internal, customer and learning perspectives?
- ☐ Is target performance set using an appropriate balance between historical, strategic, external and absolute performance targets?
- ☐ Are both performance and process methods benchmarked against similar operations and/or processes externally?
- ☐ Is benchmarking done on a regular basis and seen as an important contribution to improvement?
- ☐ Is some formal method of comparing actual and desired performance (such as the importance–performance matrix) used?
- ☐ To what extent does the operation have a predisposition towards breakthrough or continuous improvement?
- ☐ Have breakthrough improvement approaches such as business process re-engineering been evaluated?
- ☐ Are continuous improvement methods and problem-solving cycles used within the operation?
- ☐ If they are, has continuous improvement become a part of everyone's job?
- ☐ Has the Six Sigma approach to improvement been evaluated?
- ☐ Are the more common improvement techniques used to facilitate improvement within the operations?
- ☐ Does the operation show any signs of becoming a fashion victim of the latest improvement approach?
- ☐ Does the operation have a well thought through approach to managing improvement?

CASE STUDY

Ferndale Sands Conference Centre[15]

Mario Romano, the owner and General Manager of Ferndale Sands Conference Centre, had just seen an article in *The Conference Centre Journal*, and he was furious. The excellent reputation that he had worked so hard to build up over the last ten years was being threatened by one unreasonable customer and a piece of sloppy, sensationalist journalism. *'It really is unfair. Why do they let one mistake dictate the whole story? I'll tell you why, it's because they are more interested in a damning headline than they are in representing the truth.'*

Ferndale Sands Conference Centre is a conference venue of 52 rooms in Victoria State, Australia, about 20 km outside Melbourne. Established by Mario's father, initially as a hotel, it was re-launched as a conference centre four years ago and Mario was broadly pleased with the business he had attracted so far. The centre had managed to establish a presence in the fast-growing and profitable conference market. Specifically, it aimed at the 'executive retreat', rather than the 'company meeting' market. These events could be anything from one-day's duration through two weeks'. He had also negotiated deals with three higher education institutions to accommodate their Executive Education programmes. With its traditional 'Victorian' architecture, tranquil setting and excellent kitchen, Ferndale Sands offered a '*. . . supremely comfortable setting in which to work on those important decisions that will shape the future of your organization'* (Ferndale Sands brochure).

What had infuriated Marco was an article in the journal that had claimed to uncover administrative complacency and inefficiency at some of the State's conference centres (see the extract in Figure 12.16). Yet in the same edition, another piece had, generally, given a good rating to Ferndale Sands. This article had compared four conference centre facilities in and around Melbourne, and although the editorial comment had been neutral, the details included in the survey had quite clearly shown Ferndale Sands in a favourable light. Table 12.2 shows the summary of the four conference centres reviewed.

Both Mario and his front-of-house manager Robyn Wells disputed the article's rating of their administrative capabilities. However, they also were aware that administrative support was seen as being fairly important when they surveyed their clients (see Table 12.3 for Ferndale Sands' latest survey results).

Mario was determined to do something about the negative publicity. He called a meeting between himself, Nick Godfrey who was in charge of catering and recreational facilities and Robyn Wells the front-of house manager, who also was in charge of all client relations.

Nick –*'OK Mario, I know you're not pleased, but I think you are in danger of overreacting. The best way to respond is just to ignore it. It's the survey that will be saved by potential clients, not a minor article at the front of the journal. And it's the survey that reflects what we really are.'*

Shambles behind the grandeur?

Alison Peraway

Even the grandest of Victorian conference centres can fall from grace it seems. Recent complaints that Ferndale Sands may look like a Governor's palace, but can't get the basics right, were supported in hard hitting comments from the State's leading reservations agency. *"Ferndale Sands may not be the only venue to get complacent said Charles*

Figure 12.16 Extract from *The Conference Journal*

Table 12.2 Extract from *The Conference Centre Journal's* survey of conference venues in and around Melbourne

	Ferndale Sands Conference Centre	*Collins International*	*The Yarrold Conference Centre*	*St Kildan Conference Centre*
Price ($$$$$ = very expensive, $$$ = average, $ = budget)	$$$$	$$$$$	$$$	$
Size of menu	Extensive	Standard	Standard	Limited
Decor	Traditional, luxurious and tasteful,	Modern, very luxurious and stylish	Modern but basic	Needs renovation
Style and quality of food	Modern, best in the State	International modern, slightly standardised	Varies, undistinguished but acceptable	Varies, very basic
Quality of service	Excellent, friendly and relaxed	Good	Limited	Enthusiastic but limited
Administration and support	Variable, some problems recently	Good	Good	OK
Flexibility of accommodation	Poor	Very good	Acceptable	OK
Off-peak price discounts	None offered	Some in summer, none in winter	Willing to negotiate	Willing to negotiate
Equipment	Normal range	Normal range	Normal range	Requires notice for 'anything unusual'
Recreational facilities	Full range, gym, tennis, golf, swimming pool, etc.	Gym, swimming pool	Gym only	Gym only
Ease of access	Good, will pick up from airport and city	City centre, no airport shuttle to hotel, but bus service, taxies, etc.	Close to city centre	10 km from city centre

Table 12.3 Percentage of Ferndale Sands clients reporting factor as important or very important

	Percentage of customers reporting factor as important or very important (%)
Price	72
Size of menu	16
Decor	55
Style and quality of food	58
Quality of service	89
Administration and support	56
Flexibility of accommodation	85
Off-peak price discounts	16
Equipment	72
Recreational facilities	21
Ease of access	73

Mario –*'Yeh, but even that gets its wrong. It shows administration and support as uneven. What do they mean uneven?'*

Robin –*'I don't know, I guess they must have talked to a couple of clients with some kind of grudge. But look, two things always come out as the most important things for our customers: quality of service and the flexibility to accommodate their needs for different configuration of rooms. We're great at service quality. We're always getting extravagant praise; it's a real winner for us. I've got files full of compliments. It's room flexibility that's our problem. Most clients accept that you can't mess around with a historic building like this, but that doesn't get round the fact that we can't reconfigure our rooms like you could in a modern hotel with sliding room partitions.'*

Mario –*'Well maybe that's something we could minimise by making it clear to clients what we can and can't do when they make a reservation.'*

Robyn –*'True, and we do that when demand is very high. But you can't ask us to turn away business by stressing what we can't do during quiet periods.'*

Mario –*'Well maybe we should. But that's not my main concern right now. What worries me are the things that always show up as mid-range factors in our customer surveys. We tend to forget about these. They may not be the most* *important things in the clients' eyes but they are not unimportant either. I'm talking about things like the quality of our food and the décor of the rooms, and also Robyn, the administrative support we offer. If we get these things wrong, it can almost cause us as many problems as the really important things. That's why I'm upset about the poor administrative score we get in the journal. We score five for décor, and really good for food, but poor for administration.'*

Robyn –*'But as we said, that's just unfortunate. I still have every confidence in my administrative staff.'*

Nick –*'Before we get into that again, can I raise the question of our recreation facilities? It's one our best assets, yet it never rates as important with clients. It's the same with the choice we offer on our menus. Both these things are expensive to provide, and yet we don't seem to get the benefit. Why don't we make a real effort to really promote both of these things? You know, really convince the clients that great facilities and a wide choice on the menu are things that are worth paying a little more for.'*

Mario –*'It's not our pricing that's the problem. Although it's a fairly important issue with most clients, we can command relatively high prices. It's our costs that worry me more. Our general running costs are higher than they should be. Talking to the guys over at Parramatta Pacific in New South Wales, they are very similar to us, but their costs are a good 10 per cent less than ours. '*

Robyn –*'So, what is more important, raising our revenue or cutting our costs? '*

Mario –*'They are both equally important of course. The point is, what do we do about attracting more business and keeping our costs down?'*

Robyn –*'OK, we've got to do something, but remember we've also got the centre to run. Our busy period is just coming up and I don't want everyone distracted by lots of little improvement initiatives.'*

Nick –*'Absolutely. We have to limit ourselves to one or two actions that will have a noticeable impact.'*

Mario –*'I think you're probably right. But I would also add a further comment. And that is that if what we choose to do requires investment, then it must be guaranteed to have an impact. I need to go now, but why don't you two draw up a list of things that we could do. I'll review them later. OK thanks everyone.'*

Robyn and Nick's suggestions

In fact Robyn and Nick decided to draw up their individual lists of potential improvement initiatives. They also decided that, to begin with, only two of these improvement initiatives should be chosen.

Robyn's suggestions:

1 Increase prices – *'Why not? Although demand is variable, the general trend is rising as the conference market expands. At this top end of the market, I don't think we are that price sensitive. It would also bring in the revenue that we need to make further reinvestments to the centre.'*

2 Reduce menu choice – *'This really is a left-over from what menus used to be like. It goes back to the time when 'more' was considered 'better' instead of just 'more'. It is also very expensive to maintain that range of food while still maintaining quality.'*

3 Close the golf course – *'The golf course is probably the most expensive facility we have outside the house. It isn't rated by customers, so why do we keep it on?'*

4 Renovate outbuilding to provide flexible conference rooms – *'We can't easily change the inside of the house, but we do have outbuildings that could be converted into conference suites. They could be equipped with moveable partitions that would enable the space to be configured however our clients wanted it. OK, it would be expensive, but in the long-term it's necessary.'*

Nick's suggestions:

1 Promote food and facilities more effectively – *'We have a great reputation for food and for having marvellous facilities. Ferndale Sands is just a beautiful place, yet we're not exploiting it fully. A campaign from a good public relationship company could really establish us as the premier conference centre, not just in the State, but in the whole of Australia.'*

2 Cut in-house staff numbers and replace part-time staff with a smaller number of full-time staff – *'Having so many part-time staff is expensive. We pay them the same hourly rate as full-time staff, yet there are all the extra personnel costs. Also, I think we are over-staffed in the house. Staff costs are a major part of our expenditure. It's the obvious area to look for cuts.'*

3 Invest in more equipment, both relaxation equipment and presentation equipment – *'We have great sporting facilities but they could be better. If we are going to exploit them more, it may not be a bad thing to invest even more heavily in them. Also, we could make sure that we were ahead of the game in terms of the very latest audio-visual equipment. Both these things would help us to promote ourselves as the premier conference centre in Australia.'*

QUESTIONS

1 What factors would you use to judge the operations performance of Ferndale Sands'?

2 What improvement priority would you give to each of these performance measures?

3 Which two suggestions put forward by Robyn and Nick would you recommend?

APPLYING THE PRINCIPLES

Some of these exercises can be answered by reading the chapter. Others will require some general knowledge of business activity and some might require an element of investigation. **Model answers for the first two questions are available on the student companion website.** For model answers to the other questions in this section, please ask your tutor/lecturer.

1 'Everything we do can be broken down into a process', said Lucile, COO of an outsourcing business for the 'back-office' functions of a range of companies. 'It may be more straightforward in a manufacturing business, but the concept of process improvement is just as powerful in service operations. Using this approach, our team of Black Belts has achieved 30 per cent productivity improvements in 6 months. I think Six Sigma is powerful because it is the process of comparing process outputs against customer requirements. To get processes operating at less than 3.4 defects per million opportunities means that you must strive to get closer to perfection and it is the customer that defines the goal. Measuring defects per opportunity means that you can actually compare the process of, say, a human resources process with a billing and collection process.'

(a) What are the benefits of being able to compare the amount of defects in a human resources process with those of collection or billing?

(b) Why is achieving defects of less than 3.4 per million opportunity seen as important by Lucile?

(c) What do you think are the benefits and problems of training Black Belts and taking them off their present job to run the improvement projects, rather than the project being run by a member of the team that has responsibility for actually operating the process?

2 Develop cause–effect diagrams for the following types of problem:

(a) staff waiting too long for their calls to be answered at their IT helpdesk

(b) poor food in the company restaurant

(c) poor lecturing from teaching staff at a university

(d) customer complaints that the free plastic toy in their breakfast cereal packet is missing

(e) staff having to wait excessively long periods to gain access to the coffee machine.

3 For over 10 years a hotel group had been developing self-managed improvement groups within its hotels. At one hotel reception desk, staff became concerned about the amount of time the reception desk was left unattended. To investigate this, the staff began keeping track of the reasons they were spending time away from the desk and for how long each absence kept them away. Everyone knew that reception desk staff often had to leave their post to help or give service to a guest. However, no one could agree what was the main cause of absence. Collecting the information was itself not easy because the staff had to keep records without affecting customer service. After three months the data were presented in the form of a Pareto diagram, which is shown in Table 12.4. It came as a surprise to reception staff and hotel management that making photocopies for guests was the main reason for absence. Fortunately, this was easily remedied by moving the photocopier to a room adjacent to the reception area, enabling staff to keep a check on the reception desk while they were making copies.

(a) Do you think it was wise to spend so much time on examining this particular issue? Isn't it a trivial issue?

(b) Should this information be used to reflect improvement priorities? In other words, was the group correct to put priority on avoiding absence through photocopying, and should its next priority be to look at the time spent checking files in the back office?

4 A transport services company provides a whole range of services to railway operators. Its reputation for quality was a valuable asset in its increasingly competitive market. '*We are continually looking for innovation in the*

Table 12.4 Reasons for staff time away from the reception desk

Reason for being away from reception desk	*Total number of minutes away*
Checking files in back office	150
Providing glasses for night drinks	120
Providing extra key cards	90
Providing medication	20
Providing extra stationery	70
Providing misc. items to rooms	65
Providing night drinks	40
Making photocopies	300
Carrying messages to meeting rooms	125
Locking and unlocking meeting rooms	80
Providing extra linen	100

way we deliver our services because the continuous improvement of our processes is the only way to make our company more efficient', said the company's CEO. '*We use a defined set of criteria to identify critical processes, each of which is allocated a "process owner" by our quality steering committee. This is helped by the company's "process excellence index" (EPI), which is an indicator of the way a process performs, particularly how it is designed, controlled and improved. The EPI score, which is expressed on a scale of 1 to 100, is calculated by the process owner and registered with the quality department. With this one figure we can measure the cost, reliability and quality of each process so that we can compare performance. If you don't measure, you can't improve. And if you don't measure in the correct way, how can you know where you are?. Our suggestion scheme is designed to encourage staff to submit ideas that are evaluated and rated. No individual suggestion is finally evaluated until it has been fully implemented. Where a team of employees puts ideas forward, the score is divided between them, either equally or according to the wishes of the team itself. These employee policies are supported by the company's training schemes, many of which are designed to ensure all employees are customer-focused.'*

(a) What seem to be the key elements in this company's approach to improvement?

(b) Do you think this approach is appropriate for all operations?

5 (a) – As a group, identify a 'high visibility' operation that you all are familiar with. This could be a type of quick service restaurant, record stores, public transport systems, libraries, etc.

(b) – Once you have identified the broad class of operation, visit a number of them and use your experience as customers to identify:

(i) the main performance factors that are of importance to you as customers, and

(ii) how each store rates against each other in terms of their performance on these same factors.

(c) Draw an importance–performance diagram for one of the operations that indicates the priority they should be giving to improving their performance.

(d) Discuss the ways in which such an operation might improve its performance and try to discuss your findings with the staff of the operation.

Notes on chapter

1 Sources include: Robinson, D. (2013) 'Prêt A Manger to speed up expansion', *Financial Times,* 22 April; Slywotzky, A. (2012) 'Building a steep trajectory of improvement: The Prêt A Manger case', *Fast Company,* 10 October; Goodman, M. (2011) 'Pret Smile, it will pay for everyone', *Sunday Times,* 6 March; Thring, O. (2012) 'Prêt A Manger's success is deserved', The *Guardian,* 3 April; Prêt A manger website, http://www.pret.com/

2 Lewis Carroll (1871) *Alice Through the Looking Glass.*

3 Lebas, M.J. (1995) 'Performance measurement and performance management', *International Journal of Production Economics,* vol. 41, 23–35.

4 See Kaplan, R.S. and Norton, D.P. (1996) *The Balanced Scorecard,* Harvard Business School Press, Boston.

5 Kaplan, R.S. and Norton, D.P. (1995) op.cit.

6 *The Economist* (2016) 'The great escape: What other makers can learn from the revival of Triumph motorcycles', Print Edition, 23 Jan 23.

7 Shenkar, O. (2010) *Copycats: How Smart Companies Use Imitation to Gain a Strategic Edge,* Harvard Business School Press

8 Shenkar, O. (2012) 'Just imitate it! A copycat path to strategic agility', *Ivey Business Journal,* May / June.

9 Sources include: West, K. (2011) 'Formula One trains van drivers', *The Times,* 1 May; Formula One website, http://www.f1network.net/main/s107/st164086.htm

10 Ferdows, K. and de Meyer, A. (1990) 'Lasting improvement in manufacturing', *Journal of Operations Management,* vol. 9, no. 2. However, research for this model is mixed. For example, Patricia Nemetz questions the validity of the model, finding more support for the idea that the sequence of improvement is generally dictated by technological (operations resource) or market (requirements) pressures; see Nemetz, P. (2002) 'A longitudinal study of strategic choice, multiple advantage, cumulative model and order winner/qualifier view of manufacturing strategy', *Journal of Business and Management,* January.

11 Hammer, M. and Champy, J. (1993) *Re-engineering the Corporation,* Nicholas Brealey Publishing.

12 Hammer, M. (1990) 'Re-engineering work: don't automate, obliterate', *Harvard Business Review,* vol. 68, no. 4.

13 Pande, P.S., Neuman, R.P. and Cavanagh, R.R. (2000) *The Six Sigma Way,* McGrawHill, New York.

14 Sources: Gawande A (2010)The Checklist Manifesto: how to get things right, *Metropolitan*; Aaronovitch, D. (2010) 'The Checklist Manifesto: review', The Times, 23 January.

15 This case is based on several residential conference centres but does not reflect the concerns of any particular one.

TAKING IT FURTHER

Ahlstrom, J. (2014) ***How to Succeed with Continuous Improvement: A Primer for Becoming the Best in the World, McGraw-Hill Education.*** *Based on a real-life case study*

George, M. Rowlands, L., D. and Kastle, B. (2003) ***What Is Lean Six Sigma? McGraw-Hill Publishing Co.*** *Very much a quick introduction on what Lean Six Sigma is and how to use it.*

Goldratt, E. M. (2012) ***The Goal: A Process of Ongoing Improvement, Gower Publishing Limited, Updated version of a classic***

Hindo, B. (2007) ***'At 3M, A Struggle between Efficiency and Creativity: how CEO George Buckley is managing the yin and yang of Discipline and Imagination, Business week, 11 June.*** *Readable article from the popular business press.*

Pande, P.S., Neuman, R.P. and Cavanagh, R. (2002) ***Six Sigma Way Team Field Book: An Implementation Guide for Project Improvement Teams, McGraw Hill.*** *This is a unashamedly practical guide to the Six Sigma approach.*

Paper, D. J., Rodger, J A. and Pendharkar, P C. (2001) ***'A BPR case study at Honeywell', Business Process Management Journal, vol. 7, no. 2, 85–99.*** *Interesting, if somewhat academic, case study.*

Zu, X., Fredendall, L D. and Douglas, T J. (2008) ***'The evolving theory of quality management: The role of Six Sigma', Journal of Operations Management, vol. 26, 630–650.*** *As it says. . .*

13

Quality management

Introduction

No business, or enterprise of any kind, can afford to ignore quality. It can build a company's reputation, allowing it to preserve its margins, and it can save the operational and reputational costs of poor quality. That is why all businesses are concerned with quality, usually because they understand that high quality can be a significant competitive advantage. But 'quality management' has come to mean more than avoiding errors. It is an approach to the way in which to manage and more significantly, improve processes generally. This is because quality management focuses on the very fundmentals of operations and process management: the ability to produce and deliver the products and services that the market requires, both in the short and long term. A grasp of quality management principles is the foundation of any improvement activity. Figure 13.1 shows the position of the ideas described in this chapter in the general model of operations management.

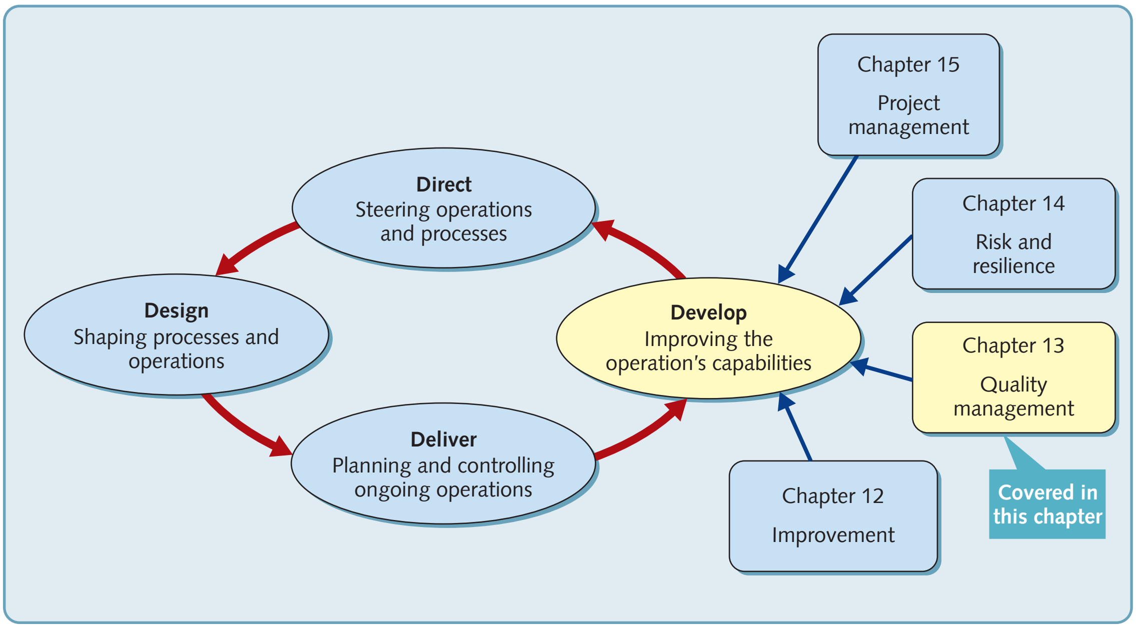

Figure 13.1 **Quality management is the activity of ensuring consistent conformance to customers' expectations**

EXECUTIVE SUMMARY

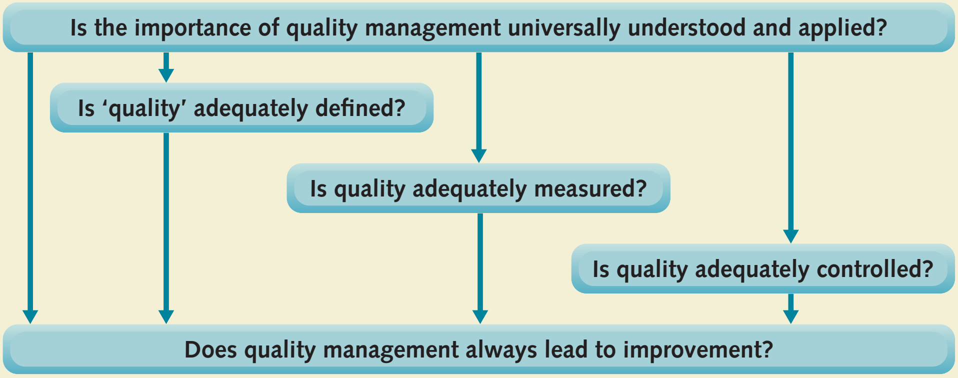

Is the importance of quality management universally understood and applied?

Quality is consistent conformance to customers' expectations. Managing quality means ensuring that an understanding of its importance and the way in which it can be improved, are spread throughout the business. There has been significant development in this subject over the last several decades but, arguably, the most recent and most significant impact on how to manage quality has come from the total quality management (TQM) movement. Quality management is now something that can be universally applied throughout a business and by implication also is the responsibility of all managers in the business. In particular, it applies to all parts of the organisation. The internal customer concept can be used to establish the idea that it is important to deliver a high-quality service to internal customers (other processes in the business). Service level agreements can be used to operationalise the internal customer concept. Just as important, is the idea that quality also applies to every individual in the business. Everyone has the ability to impair quality, so everyone also has the ability to improve it.

Is quality adequately defined?

Quality needs to be understood from the customer's point of view because it is defined by the customer's perceptions and expectations. One way of doing this is to use a quality gap model. This starts from the fundamental potential gap between customers' expectations and perceptions and deconstructs the various influences on perceptions and expectations. Gaps between these factors are then used to diagnose possible root causes of quality problems. A further development is to define the quality characteristics of products or services in terms of their functionality, appearance, reliability, durability, recovery and contact.

Is quality adequately measured?

Without measuring quality it is difficult to control it. In addition, the various attributes of quality can be measured either as a variable (measured on a continuously variable scale), or attribute (a binary acceptable or not acceptable judgement). One approach to measuring quality is to express all quality-related issues in cost terms. Quality costs are usually categorised as prevention costs (incurred in trying to prevent errors), appraisal costs (associated with checking for errors), internal failure costs (errors that are corrected with the operation) and external failure costs (errors that experienced by customers). Generally, increasing expenditure on prevention will bring a more than equivalent reduction in other quality-related costs.

Is quality adequately controlled?

Control means monitoring and responding to any deviations from acceptable levels of quality. One of the most common ways of doing this is through statistical process control (SPC). This technique not only attempts to reduce the variation in quality performance, to enhance process knowledge, but also is used to detect deviations outside the 'normal' range of quality variation.

Does quality management always lead to improvement?

Very often quality improvements do not sustain because there is no set of systems and procedures to support and embed them within the operation's day-to-day routines. The best-known system for doing this is the ISO 9000 approach adopted now throughout the world. Of the other systems, one of the most widely known is the European Foundation for Quality Management (EFQM) excellence model. Once only known as the basis of the European Quality Award, it is now extensively used as a self-assessment tool that allows organisations to assess their own quality systems.

DIAGNOSTIC QUESTION

Is the importance of quality management universally understood and applied?

Before looking at the 'how' and 'why' of quality management, it is useful to understand what the term can be taken to mean.

What is quality management?

There are many definitions of 'quality': 'conformance to specification', being 'fit for purpose', 'achieving appropriate specification', and so on. The one we use here is. . . *'quality is consistent conformance to customers' expectations'* because it includes both the idea of quality as *specification* (what the product or service can do), and quality as *conformance* (there are no errors, so it always does what it is supposed to do). Not surprisingly, for such an important topic, it has a history. Approaches to quality management have always been of interest to any business that aspired to satisfy its customers. Arguably, the most significant of the approaches to quality management was total quality management (TQM) that became popular with all types of business in the late 1970s and 1980s. Although it was based on earlier work by several management thinkers, Feigenbaum popularised the term 'total quality management' in 1957. After that, it was developed through the work of several 'quality gurus' including Deming, Juran, Ishikawa, Taguchi and Crosby (see 'Taking it further' at the end of the chapter).

It can be viewed as a logical extension of the way in which quality-related practice has progressed. Originally, quality was achieved by inspection – screening out defects before customers noticed them. Then the 'quality control' (QC) concept developed a more systematic approach to not only detecting, but also solving quality problems. 'Quality assurance' (QA) widened the responsibility for quality to include functions other than direct operations, such as human resources, accounting and marketing. It also made increasing use of more sophisticated statistical quality techniques. TQM included much of what went before but developed its own distinctive themes, especially in its adoption of a more 'all embracing' approach. Since the fashionable peak of TQM, there has been some decline in the status of TQM, yet its ideas, many of which are included in this chapter, have become accepted quality practice. The two businesses described in the following examples both incorporate TQM ideas in their approach to quality, especially their inclusion of every employee.

EXAMPLE

The Swiss Army knife: 'Our best means of protection is quality'[1]

The famous Swiss Army Knife is known all over the world for its usefulness and quality. It is made by the Victorinox Company in its factory in the small Swiss town of Ibach, Canton Schwyz and has a history that can be traced back to 1891. The company has numerous letters from its customers testifying to their product's quality and durability. The following story from one engineer is typical. '*I was installing a new piece of equipment in a sewage treatment plant. One morning, as I was crossing the bridge over the aeration tank of the treatment plant, I saw that the setting on one of the instruments was incorrect. I took out my Swiss Army knife to make the necessary adjustment. The knife slipped out of my hand and fell into the aeration tank, whose function is to oxidise organic waste – the oxidising environment that is extremely corrosive to metals. Four years later, I received a small parcel with a note from the supervisor*

of the plant. They had emptied the aeration tank and found my knife at the bottom. The parcel contained the knife that was in astonishingly good condition. The plastic casing and cover had only suffered very minor damage. I can assure you that very few products could have survived treatment like this, the components would have dissolved or simply disappeared.'

Today, the Victorinox factory assembles 27,000 knives a day (plus nearly 100,000 other items). More than 450 steps are required in its manufacture. However, times have not been easy for the Victorinox Company. Airport security restrictions after 9/11 hit sales of the knife. '*Our sales plummeted almost overnight*', said Carl Elsener, the Company's CEO and the great-grandson of its founder. '*All airport shops were suddenly banned from selling knives and we lost 30 per cent of our income that came from spontaneous airport purchases.*' But rather than shut down some of its production lines and get rid of a considerable chunk of its workforce to cut costs (the factory hasn't fired a single person for economic reasons in all of the125 years of its existence). Victorinox developed new products including laser-fronted ballpoint pens, bladeless 'in-flight' knives, Swiss Memory and Swiss Flash foldable USB drives. Another major threat to sales that has been growing is the appearance on the market of fake 'Swiss Army' knives, made mostly in China. Many of them look similar to the original; they even have the familiar Swiss cross on the handle.

So what is their defence against these fakes? '*Quality*', said Carl Elsener. '*We have exhausted all legal means for the brand protection of our popular products. Our best means of protection is quality which remains unsurpassed and speaks louder than words.*' And the three components of the 'Victorinox quality control system' are at the heart of this quality defence. First, incoming materials are checked to conform to quality specifications. Non-conforming products are identified, evaluated and reviewed according to set procedures. Only steel and plastic that complies with Victorinox rigorous quality standards are used in the manufacture of the products. Second, process control is employed at all stages of the production process. Third, the Final Inspection Department employs 50–60 people who are responsible for ensuring that all products conform to requirements. Any non-conforming products are isolated and identified. Non-conforming parts are repaired or replaced at the repair department.

EXAMPLE

The Four Seasons Hotel Canary Wharf[2]

The first Four Seasons Hotel opened over 45 years ago. Since then the company has grown to 81 properties in 34 countries. Famed for its quality of service, the hotel group has won countless awards, including the prestigious Zagat survey, and numerous AAA Five Diamond Awards. And it is also one of only 14 organisations that have been on the FORTUNE magazine's list of '100 Best Companies to Work For' every year since it launched in 1998, thus ranking as 'top hotel chain' internationally. From its inception, the group has had the same guiding principle, 'to make the quality of our service our competitive advantage'. The company has what it calls its Golden Rule: 'Do to others (guests and staff) as you would wish others to do to you'. It is a simple rule, but it guides the whole organisation's approach to quality.

'Quality service is our distinguishing edge and the company continues to evolve in that direction. We are always looking for better, more creative and innovative ways of serving our guests', said Michael Purtill, the general manager of the Four Seasons Hotel Canary Wharf in London. *'We have recently refined all of our operating standards across the company, enabling us to further enhance the personalised, intuitive service that all our guests receive. All employees are empowered to use their creativity and judgement in delivering exceptional service and making their own decisions to enhance our guests' stay. For example, one morning an employee noticed that a guest had a flat tyre on their car and decided on his own accord to change it for them, which was very much appreciated by the guest.'*

'The Golden Rule means that we treat our employees with dignity, respect and appreciation. This approach encourages them to be equally sensitive to our guests' needs and offer sincere and genuine service that exceeds expectations. Just recently, one of our employees accompanied a guest to the hospital and stayed there with him for an entire afternoon. He wanted to ensure that the guest wasn't alone and was given the medical attention he needed. The following day that same employee took the initiative to return to the hospital (even though it was his day off) to visit and made sure that that guest's family in America was kept informed about his progress.'

'We ensure that we have an ongoing focus on recognising these successes and publicly praise and celebrate all individuals who deliver these warm, spontaneous, thoughtful touches.'

'At Four Seasons we believe that our greatest asset and strength are our people. We pay a great deal of attention to selecting the right people with an attitude that takes great pride in delivering exceptional service. We know that motivated and happy employees are essential to our service culture and are committed to developing our employees to their highest potential. Our extensive training programs and career development plans are designed with care and attention to support the individual needs of our employees as well as operational and business demands. In conjunction to traditional classroom based learning, we offer tailor-made internet based learning featuring exceptional quality courses for all levels of employee. Such importance is given to learning and development that the hotel has created two specialised rooms, designated to learning and development. One is intended for group learning and the other is equipped with private computer stations for internet-based individual learning. There is also a library equipped with a broad variety of hospitality related books, CDs and DVDs, which can be taken home at any time. This encourages our employees to learn and develop at an individual pace. This is very motivating for our employees and in the same instance, their development is invaluable to the growth of our company. Career wise, the sky is the limit and our goal is to build lifelong, international careers with Four Seasons.'

'Our objective is to exceed guest expectations and feedback from our guests and our employees is an invaluable barometer of our performance. We have created an in-house database that is used to record all guest feedback (whether positive or negative). We also use an online guest survey and guest comment cards, which are all personally responded to and analysed to identify any potential service gaps.'

'We continue to focus on delivering individual personalised experiences and our Guest History database remains vital in helping us to achieve this. All preferences and specific comments about service experience are logged on the database. Every comment and every preference is discussed and planned for, for every guest and for every visit. It is our culture that sets Four Seasons apart; the drive to deliver the best service in the industry that keeps their guests returning again and again.'

What do they have in common?

Both companies have built their reputation on providing exceptional quality; Victorinox by producing products that are both durable and functional, Four Seasons by producing the type of immaculate service that its guests have come to expect. Just as important, both companies understand that their quality depends on rigorously controlling the processes through which their products and services are produced. Customers are willing to pay more for a 'Swiss Army Knife' because they trust Victorinox to produce consistently good products. The guests at the Four Seasons are paying for, and expect, exceptional service at a top-range hotel. Both Victorinox and Four Seasons try to conform to what they promise, even though exactly what they promise is expressed in different ways. Both businesses define quality in terms of the expectations and perceptions of customers. This means seeing things from a customer's point of view. Customers are not regarded as *external* to the organisation but as the most important *part* of it. Both see quality, not as a single attribute, but a combination of many different things, some of which are difficult to define. There is also an emphasis on every part of the business and every individual having responsibility for ensuring quality.

OPERATIONS PRINCIPLE

Quality is multi-faceted, its individual elements differ for different operations.

Quality management should be universally understood and applied

If an operation is to fully understand customers' expectations and to match or exceed them in a consistent manner, it needs to take a universal or *total* approach to quality. Adopting a universal approach means that an understanding of *why* quality is important and *how* quality can be improved permeates the entire organisation. This idea was popularised by proponents of total quality management (TQM), who saw TQM as the ideal unifying philosophy that could unite the whole business behind customer-focused improvement. In particular, two questions are worth asking. First, does quality apply to all parts of the organisation? And second, does every person in the organisation contribute to quality?

Why is quality so important?

Put simply, quality is important because it has such a significant impact on both the revenues and costs of any business. Figure 13.2 illustrates the various ways in which quality improvements can affect other aspects of operations performance. Revenues can be increased by better sales and enhanced prices in the market. At the same time, costs can be brought down by improved efficiencies, productivity and the use of capital. A key task of the operations function therefore must be to ensure that it provides quality goods and services, both to its internal and external customers.

Does quality apply to all parts of the organisation?

If quality management is to be effective, every process, must work properly together. This is because every process affects and in turn is affected by others. Called the *internal customer concept,* it is recognition that every part of an organisation is both an internal customer and, at the same time, an internal supplier for other parts of the organisation. This means that errors in the service provided within an organisation will eventually affect the product or service that reaches the external customer. So, one of the best ways of satisfying external customers is to satisfy internal customers. This means that each process has a responsibility to manage its own internal customer–supplier relationships by clearly defining their own and their customers' exact requirements. In fact, the exercise replicates what should be going on for the whole operation and its external customers.

OPERATIONS PRINCIPLE

An appreciation of, involvement in, and commitment to quality should permeate the entire organisation.

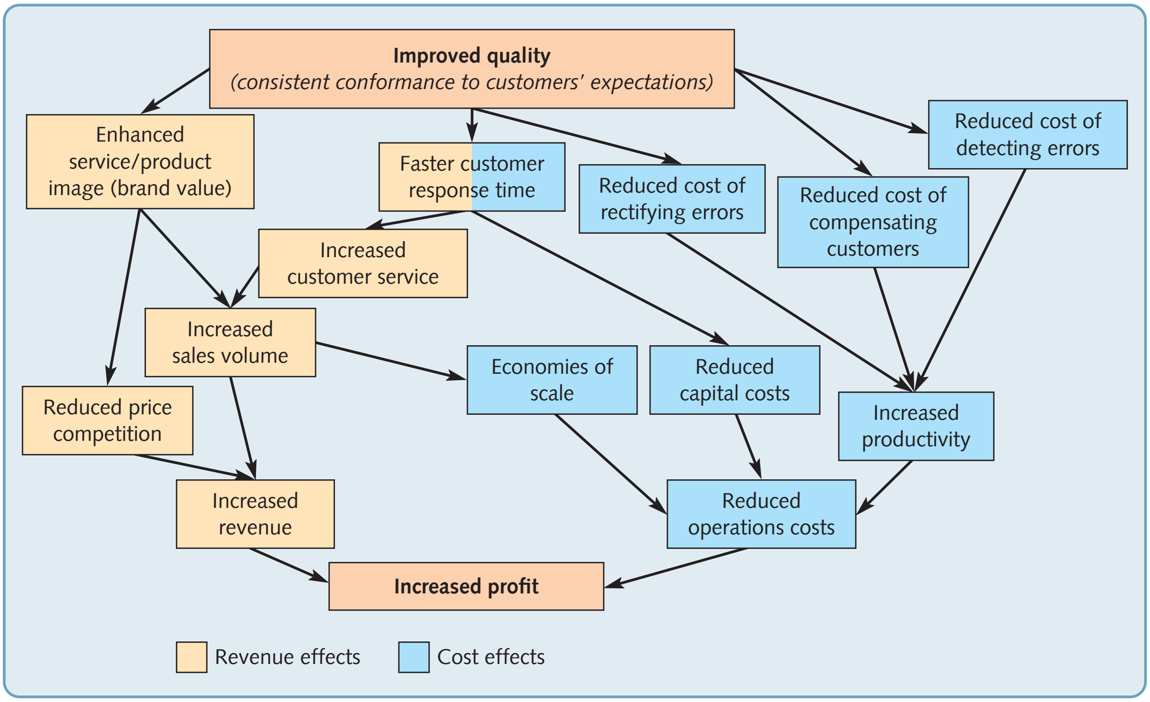

Figure 13.2 Higher quality has a beneficial effect on both revenues and costs

Service level agreements

Some operations bring a degree of formality to the internal customer concept by requiring processes to agree service level agreements (SLAs) with each other. SLAs are formal definitions of the service and the relationship between two processes. The type of issues that would be covered by such an agreement could include response times, the range of services, dependability of service supply, and so on. Boundaries of responsibility and appropriate performance measures could also be agreed. For example, an SLA between an information systems help desk and the processes that are its internal customers could define such performance measures as:

- the types of information network services that may be provided as 'standard'
- the range of special information services that may be available at different periods of the day
- the minimum 'up time', i.e. the proportion of time the system will be available at different periods of the day
- the maximum response time and average response time to get the system fully operational should it fail
- the maximum response time to provide 'special' services, and so on.

SLAs are best thought of as an approach to deciding service priorities between processes, and as a basis for improving process performance from the internal customers' perspective. At their best, they can be the mechanism for clarifying exactly how processes can contribute to the operations as a whole. See the critical commentary at the end of the chapter for a more cynical view.

Operating / operational level agreements (OLAs)

Often confused with service level agreements, an operating (or operational as they can sometimes be called, confusingly) level agreement (OLA) is an agreement that defines how various groups within an operation propose to work together to deliver what has been specified in the service level agreement. The difference between a service level agreement and an operational level agreement lies in where the boundary is drawn between the internal provider of a service and the internal customer. An SLA is the service that an internal customer should receive; an OLA is how the parts of the operation that produce the service will go about producing the service in practice. Originally, the idea of an OLA was designed to overcome the tendency of information technology (IT) functions to work in silos rather than agreeing the specific set of IT services that each part of the IT function was responsible for. The term is still more common in IT services than in other parts of a business, but is becoming a general term.

Does every person in the organisation contribute to quality?

A total approach to quality should include every individual in the business.

People are the source of both good and bad quality and it is everyone's personal responsibility to get quality right. This applies not only to those people who can affect quality directly and have the capability to make mistakes that are immediately obvious to customers, for example those who serve customers face to face or physically make products. It also applies to those who are less directly involved in producing products and services. The keyboard operator who keyboards data incorrectly, or the product designer who fails to investigate thoroughly the conditions under which products will be used in practice, could also set in motion a chain of events that customers eventually see as poor quality.

It follows that if everyone has the ability to impair quality, they also have the ability to improve it – if only by 'not making mistakes'. But their contribution is expected to go beyond a commitment not to make mistakes; they are expected to bring something positive to the way they perform their jobs. Everyone is capable of improving the way in which they do their own jobs and practically everyone is capable of helping others in the organisation to improve theirs. Neglecting the potential inherent in all people is neglecting a powerful source of improvement.

EXAMPLE

Fat fingers syndrome[3]

Feeling sleepy one day, a German bank worker briefly fell asleep on his keyboard when processing a €64 debit (withdrawal) from a pensioner's account, repeatedly pressing the number 2. The result was that the pensioner's account had €222 million withdrawn from it instead of the intended €64. Fortunately, the bank spotted the error before too much damage was done (and before the account-holder noticed). More seriously, the supervisor who should have checked his junior colleague's work was sacked for failing to notice the blunder (unfairly, a German labour tribunal later ruled). It is known as 'fat finger syndrome' – used to describe a person who makes keyboarding errors when chatting, tired or overstressed. For some people, like traders working in fast-moving electronic financial markets, if they press the wrong button on their keyboard, it means a potential fortune could be lost.

Fat finger trading mistakes are not uncommon. In 2009, Swiss bank UBS mistakenly ordered 3 trillion yen (instead of 30 million yen) of bonds in a Japanese video games firm. In 2005, a Japanese trader tried to sell one share of a recruitment company at 610,000 yen per share. But he accidentally sold 610,000 shares at one yen each, despite this being 41 times the number of shares available. Unlike the German example, the error was not noticed and the Tokyo Stock Exchange processed the order. It resulted in Mizuho Securities losing 27 billion yen. The head of the Exchange later resigned. What is believed to be the biggest fat finger error on record occurred in 2014, when share trades worth more than the size of Sweden's economy had to be cancelled in Tokyo. The error briefly sparked panic after a trader accidentally entered a trade worth nearly 68 trillion yen in several of the Asian country's largest blue-chip companies. The Japan Securities Dealers Association said the trader had in error put together the volume and price of a series of transactions instead of the volume alone. However, the transactions were cancelled 17 minutes after they were made and no permanent (financial) damage was done.

DIAGNOSTIC QUESTION

Is quality adequately defined?

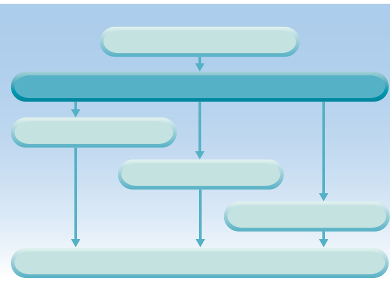

Quality is consistent conformance to customers' expectations. It needs to be understood from a customer's point of view because, to the customer, the quality of a particular product or service is whatever he or she perceives it to be. However, an individual customer's expectations may be different. Past experiences, individual knowledge and history will all shape a customer's individual expectations. Perceptions are not absolute. Exactly the same product or service may be perceived in different ways by different customers. In addition, in some situations, customers may be unable to judge the 'technical' specification of the service or product. They may then use surrogate measures as a basis for their perception of quality. For example, after seeking financial advice from an adviser it might be difficult immediately to evaluate the technical quality of the advice, especially if no better solution presents itself. In reality, a judgement of the quality of the advice may be based on perceptions of trustworthiness, relationship, the information that was provided, or the way in which it was provided.

> **OPERATIONS PRINCIPLE**
> ***Perceived quality is governed by the magnitude and direction of the gap between customers' expectations and their perceptions of a product or service.***

Closing the gaps – alignment in quality

If the product or service experience was better than expected then the customer is satisfied and quality is perceived to be high. If the product or service was less than his or her expectations then quality is low and the customer may be dissatisfied. If the product or service matches expectations then the perceived quality of the product or service is seen to be acceptable. So:

When expectations $>$ Perceptions $\rightarrow$ perceived quality is poor

When expectations $=$ Perceptions $\rightarrow$ perceived quality is acceptable

When expectations $<$ Perceptions $\rightarrow$ perceived quality is poor

Both customers' expectations and perceptions are influenced by a number of factors, some of which cannot be controlled by the operation and some of which can, at least to a certain extent. Figure 13.3 shows some of the factors that will influence the gap between expectations and perceptions and the potential gaps between some of these factors. This approach to defining quality is called a 'gap model of quality'. The model shown in Figure 13.3 is adapted from one developed by Zeithaml, Berry and Parasuraman,[4] primarily to understand how quality in service operations can be managed and identify some of the problems in so doing. However, this approach is now also used in all types of operation.

Diagnosing quality problems

Describing perceived quality in this way, allows a diagnosis of quality problems. If the perceived quality gap is such that customers' perceptions of the product or service fail to match their expectations of it, then the reason (or reasons) must lie in other gaps elsewhere in the model. Four other gaps could explain a perceived quality gap between customers' perceptions and expectations.

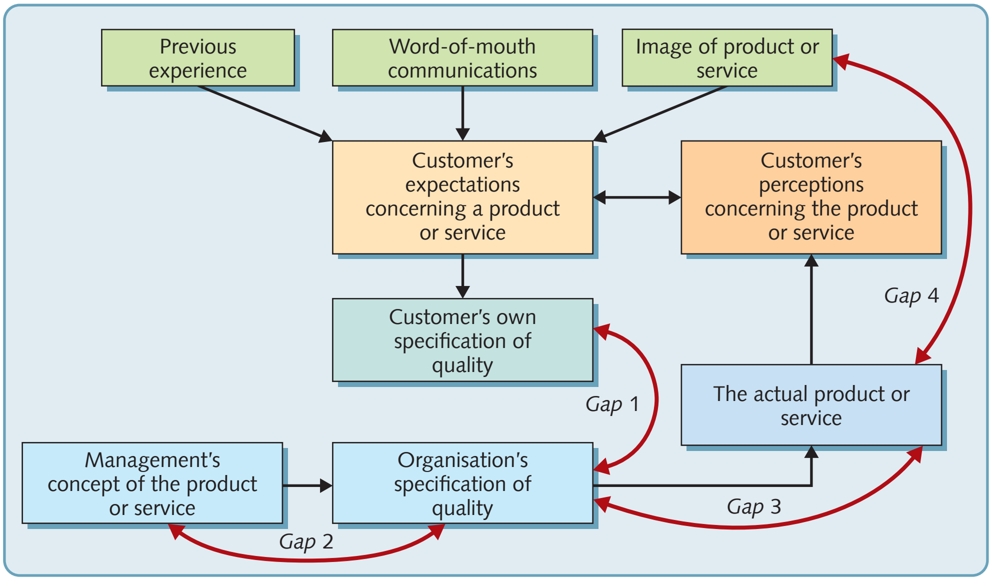

Figure 13.3 A perception–expectation gap model of quality

Gap 1: The customer's specification–operation's specification gap

Perceived quality could be poor because there may be a mismatch between the organisation's own internal quality specification and the specification that is expected by the customer. For example, a car may be designed to need servicing every 10,000 kilometres but the customer may expect 15,000 kilometre service intervals. An airline may have a policy of charging for drinks during the flight whereas the customer's expectation may be that the drinks would be free.

Gap 2: The concept–specification gap

Perceived quality could be poor because there is a mismatch between the product or service concept and the way the organisation has specified the quality of the product or service internally. For example, the concept of a car might have been for an inexpensive, energy-efficient means of transportation, but the inclusion of a catalytic converter may have both added to its cost and made it less energy efficient.

Gap 3: The quality specification–actual quality gap

Perceived quality could be poor because there is a mismatch between the actual quality of the service or product provided by the operation and its internal quality specification. This may be the result, for example, of an inappropriate or unachievable specification, or of poorly trained or inexperienced personnel, or because effective control systems are not in place to ensure the provision of defined levels of quality. For example if, despite an airline's policy of charging for drinks, some flight crews provide free drinks, they add unexpected costs to the airline and influence customers' expectations for the next flight, when they may be disappointed.

Gap 4: The actual quality–communicated image gap

Perceived quality could also be poor because there is a gap between the organisation's external communications or market image and the actual quality of the service or product delivered to the customer. This may be either the result of market positioning setting unachievable expectations in the minds of customers, or operations not providing the level of quality expected by the customer. The advertising campaign for an airline might show a cabin attendant offering to replace a customer's shirt on which food or drink has been spilt, but this service may not always be available.

EXAMPLE

Tea and Sympathy[5]

Defining quality in terms of perception and expectation can sometimes reveal some surprising results. For example, Tea and Sympathy is a British restaurant and café in the heart of New York's West Village. Over the last ten years, it has become a fashionable landmark in a city with one of the broadest ranges of restaurants in the world. Yet it is tiny, around a dozen tables packed into an area little bigger than the average British sitting room. Not only expatriate Brits, but also native New Yorkers and celebrities queue to get in. As the only British restaurant in New York, it has a novelty factor, but also it has become famous for the unusual nature of its service. *'Everyone is treated in the same way'*, said Nicky Perry, one of the two ex-Londoners who run it. *'We have a firm policy that we don't take any shit.'* This robust attitude to the treatment of customers is reinforced by 'Nicky's Rules', which are printed on the menu.

1. Be pleasant to the waitresses – remember Tea and Sympathy girls are always right.
2. You will have to wait outside the restaurant until your entire party is present: no exceptions.
3. Occasionally, you may be asked to change tables so that we can accommodate all of you.
4. If we don't need the table you may stay all day, but if people are waiting it's time to naff off.
5. These rules are strictly enforced. Any argument will incur Nicky's wrath. You have been warned.

Most of the waitresses are also British and enforce Nicky's Rules strictly. If customers object, they are thrown out. Nicky says that she has had to train 'her girls' to toughen up. *'I've taught them that when people cross the line they can tear their throats out as far as I'm concerned. What we've discovered over the years is that if you are really sweet, people see it as a weakness.'* People get thrown out of the restaurant about twice a week and yet customers still queue for the genuine shepherd's pie, a real cup of tea and, of course, the service.

Quality characteristics

Much of the 'quality' of a product or service will have been specified in its design. But not all the design details are useful in defining quality. Rather it is the *consequences* of the design that are perceived by customers. These consequences of the design are called *quality characteristics.* Table 13.1 shows a list of quality characteristics that are generally useful as applied to a service (flight) and a product (car).

Table 13.1 Quality characteristics for a motor car and an air journey

Quality characteristics	*Car*	*Flight*
Functionality – how well the product or service does its job, including its performance and features	Speed, acceleration, fuel consumption, ride quality, road-holding, etc.	Safety and duration of journey, onboard meals and drinks, car and hotel booking services
Appearance – the sensory characteristics of the product or service: its aesthetic appeal, look, feel, sound and smell	Aesthetics, shape, finish, door gaps, etc.	Décor and cleanliness of aircraft, lounges and crew

(*Continued*)

Table 13.1 Quality characteristics for a motor car and an air journey (*Continued*)

Quality characteristics	*Car*	*Flight*
Reliability – the consistency of the product's or service's performance over time, or the average time for which it performs within its tolerated band of performance	Mean time to failure	Keeping to the published flight times
Durability – the total useful life of the product or service, assuming occasional repair or modification	Useful life (with repair)	Keeping up with trends in the industry
Recovery – the ease with which problems with the product or service can be rectified or resolved	Ease of repair	Resolution of service failures
Contact – the nature of the person-to-person contact which might take place. It could include the courtesy, empathy, sensitivity and knowledge of contact staff	Knowledge and courtesy of sales and service staff	Knowledge, courtesy and sensitivity of airline staff

DIAGNOSTIC QUESTION

Is quality adequately measured?

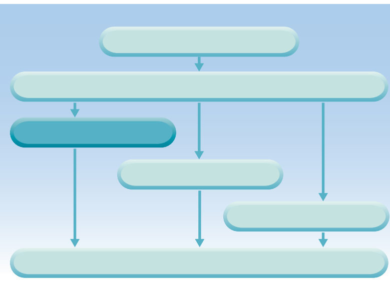

Some quality characteristics are relatively easy to measure. For example, is the gap between a car door and pillar less than 5 mm? Other more difficult-to-measure quality characteristics, such as 'appearance', need to be decomposed into their constituent elements such as 'colour match', 'surface finish' and 'the number of visible scratches', all of which are capable of being measured in a relatively objective manner. They may even be quantifiable. However, decomposing quality characteristics into their measurable sub-components can result in some loss of meaning. A quantified list of 'colour match', the 'smoothness' of the surface finish and the 'number of visible scratches' does not cover factors such as 'aesthetics', a characteristic that is difficult to measure, but nonetheless important. Some quality characteristics cannot be measured at all. The 'courtesy' of airline staff, for example, has no objective quantified measure, yet airlines place a great deal of importance on the need to ensure courtesy in their staff. In cases like this, the operation will have to attempt to measure customer *perceptions* of courtesy.

Variables and attributes

The measures used to describe quality characteristics are of two types: variables and attributes. Variable measures are those that can be measured on a continuously variable scale (for example, length, diameter, weight or time). Attributes are those that are assessed by judgement and have two states (for example, right or wrong; works or does not work; looks OK or not OK). Table 13.2 categorises some of the measures that might be used for the quality characteristics of the car and the flight.

Measuring the 'costs of quality'

One approach to measuring aggregated quality is to express all quality-related issues in cost terms. This is the 'cost of quality' approach (usually taken to refer to both costs and benefits of quality). These costs of quality are usually categorised as *prevention costs, appraisal costs, internal failure costs* and *external failure costs.* Table 13.3 illustrates the type of factors that are included in these categories.

Table 13.2 Variable and attribute measures for quality characteristics

	Car		*Flight*	
Characteristic	*Variable*	*Attribute*	*Variable*	*Attribute*
Functionality	Acceleration and braking characteristics from test bed	Is the ridge quality satisfactory?	Number of journeys which actually arrived at the destination (i.e. didn't crash!)	Was the food acceptable?
Appearance	Number of blemishes visible on car	Is the colour to specification?	Number of seats not cleaned satisfactorily	Is the crew dressed smartly?
Reliability	Average time between faults	Is the reliability satisfactory?	Proportion of journeys that arrived on time	Were there any complaints?
Durability	Life of the car	Is the useful life as predicted?	Number of times service innovations lagged competitors	Generally, is the airline updating its services in a satisfactory manner?
Recovery	Time from fault discovered to fault repaired	Is the serviceability of the car acceptable?	Proportion of service failures resolved satisfactorily	Do customers feel that staff deal satisfactorily with complaints?
Contact	Level of help provided by sales staff (1 to scale)	Did customers feel well served (yes or no)?	The extent to which customers feel well treated by staff (1 to 5 scale)	Did customers feel that the staff were helpful (yes or no)?

Table 13.3 The categories of quality cost

Category of quality-related cost	*Includes such things as. . .*
Prevention costs – those costs incurred in trying to prevent problems, failures and errors from occurring in the first place	• Identifying potential problems and putting the process right before poor quality occurs • Designing and improving the design of products and services and processes to reduce quality problems • Training and development of personnel in the best way to perform their jobs • Process control
Appraisal costs – those costs associated with controlling quality to check to see if problems or errors have occurred during and after the creation of the product or service	• The setting up of statistical acceptance sampling plans • The time and effort required to inspect inputs, processes and outputs • Obtaining processing inspection and test data • Investigating quality problems and providing quality reports • Conducting customer surveys and quality audits
Internal failure costs – failure costs that are associated with errors dealt with inside the operation	• The cost of scrapped parts and materials • Reworked parts and materials • The lost production time as a result of coping with errors • Lack of concentration due to time spent troubleshooting rather than improvement
External failure costs – failure costs that are associated with errors being experienced by customers	• Loss of customer goodwill affecting future business • Aggrieved customers who may take up time • Litigation (or payments to avoid litigation) • Guarantee and warranty costs • The cost to the company of providing excessive capability (too much coffee in the pack and too much information to a client)

Understand the relationship between quality costs[6]

At one time, it was assumed that failure costs reduce as the money spent on appraisal and prevention increases. There must be a point beyond which the cost of improving quality gets larger than the benefits that it brings. Therefore, there must be an optimum amount of quality effort to be applied in any situation that minimises the total costs of quality. Figure 13.4(a) sums up this idea.

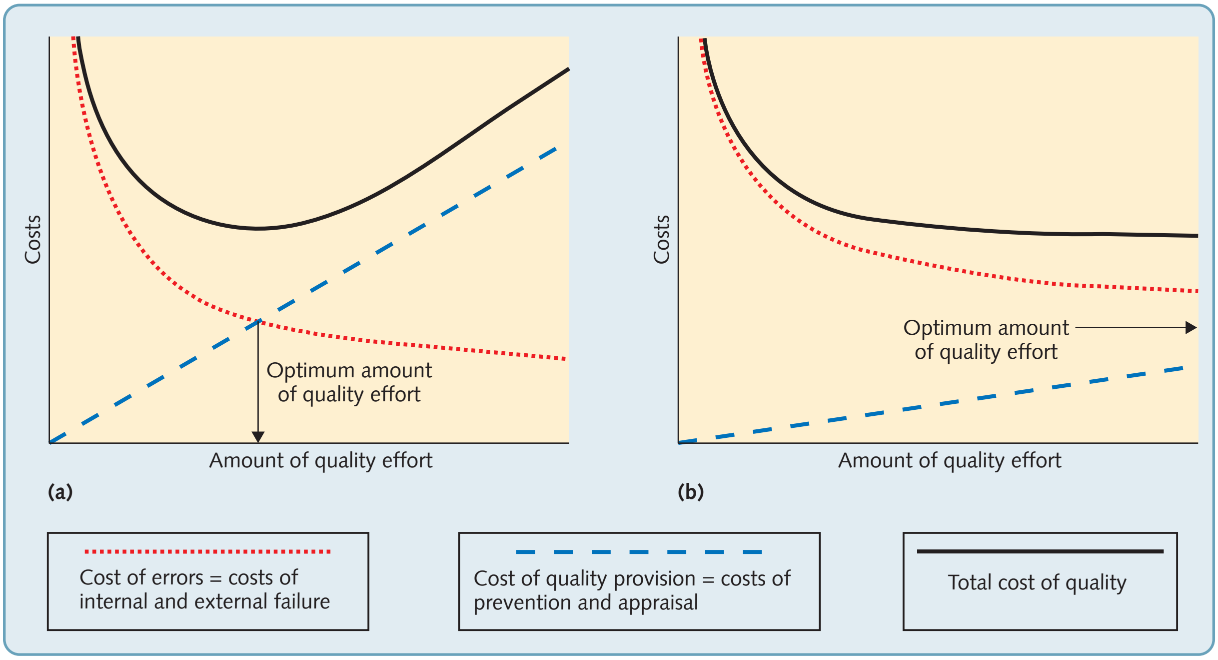

Figure 13.4 **(a) The traditional cost of quality model and (b) a more modern view**

More recently, the 'optimum-quality effort' approach has been challenged. First, why should any operation accept the *inevitability* of errors? Some occupations seem to be able to accept a zero-defect standard (even if they do not always achieve it). No one accepts the inevitability of pilots crashing a certain proportion of their aircraft, or nurses dropping a certain number of babies. Second, failure costs are generally underestimated. They are usually taken to include the cost of 'reworking' defective products, 're-serving' customers, scrapping parts and materials, the loss of goodwill, warranty costs, and so on. These are important, but in practice, the real cost of poor quality should include all the management time wasted in organising rework and rectification and, more important, the loss of concentration and the erosion of confidence between processes within the operation. Third, it implies that prevention costs are inevitably high. But by stressing the importance of quality to every individual, preventing errors becomes an integral part of everyone's work. More quality is not only achieved by using more inspectors, we all have a responsibility for our own quality and all should be capable of 'doing things right first time'. This may incur some costs – training, automatic checks, anything that helps to prevent errors occurring in the first place – but not such a steeply inclined cost curve as in the 'optimum-quality' theory. Finally the 'optimum-quality level' approach, by accepting compromise, does little to challenge operations managers and staff to find ways of improving quality.

Put these corrections into the optimum-quality effort calculation and the picture looks very different (see Figure 13.4(b)). If there is an 'optimum', it is a lot further to the right, in the direction of putting more effort (but not necessarily cost) into quality.

The TQM-influenced quality cost model

TQM rejected the optimum-quality level concept. Rather, it concentrated on how to reduce all known and unknown failure costs. Rather than placing most emphasis on appraisal (so that 'bad products and service don't get through to the customer'), it emphasised prevention (to stop errors happening in the first place). This has a significant, positive effect on internal failure costs, followed by reductions both in external failure costs and, once confidence has been firmly established, in appraisal costs. Eventually, even prevention costs can be stepped down in absolute terms, though prevention remains a significant cost in relative terms. Figure 13.5

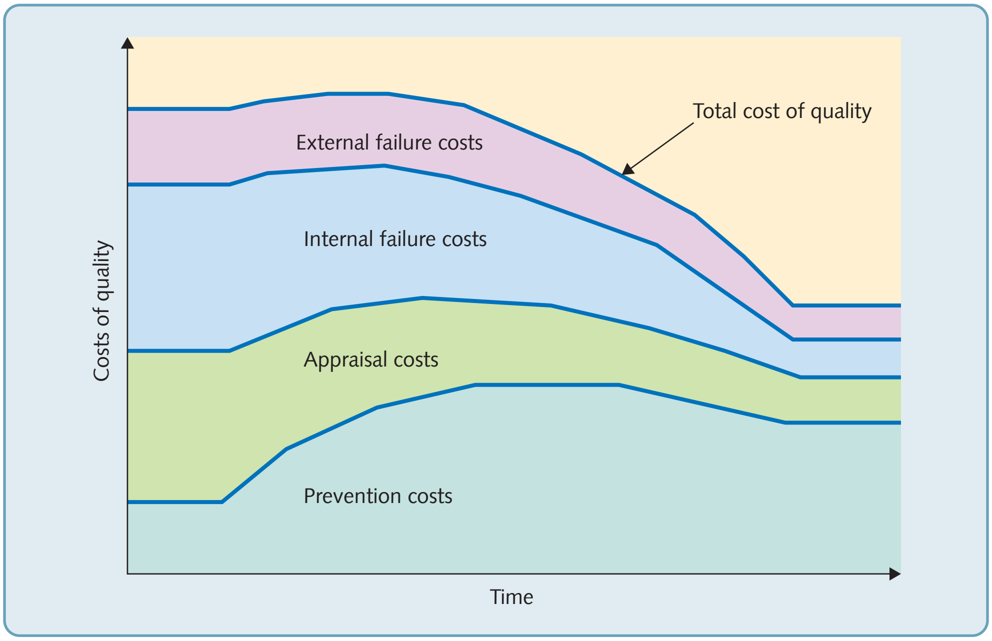

Figure 13.5 **Increasing the effort spent on preventing errors occurring in the first place brings a more than equivalent reduction in other cost categories**

OPERATIONS PRINCIPLE

Effective investment in preventing quality errors can significantly reduce appraisal and failure costs.

illustrates this idea, showing how initially total quality costs may rise as investment in some aspects of prevention is increased. Once this relationship between categories of quality cost is accepted, it shifts the emphasis from a reactive approach to quality (waiting for errors to happen, then screening them out), to a more proactive, 'getting it right first time' approach (doing something before errors happen).

DIAGNOSTIC QUESTION

Is quality adequately controlled?

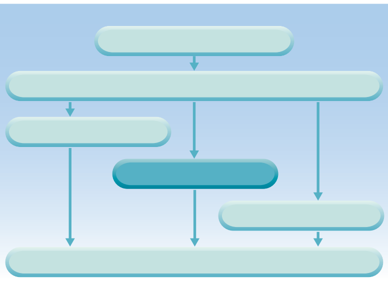

After quality has been defined and measured, processes will need to check that their quality conforms to whatever quality standards are deemed appropriate. This does not necessarily mean checking everything; sampling may be more appropriate.

Check every product and service or take a sample?

There are several reasons why checking everything may not be sensible:

- It could be dangerous to check everything. A doctor, for example, checks just a small sample of blood rather than all of it. The characteristics of this sample are taken to represent those of the rest of the patient's blood.
- Checking everything might destroy the product or interfere with the service. A lamp manufacturer cannot check the life of every single light bulb leaving the factory; they would all be

destroyed. Nor would it be appropriate for a head waiter to check every 30 seconds whether customers are enjoying the meal.
- Checking everything may be too costly. For example, it just might not be feasible to check every single item from a high-volume plastic moulding machine, or to check the feelings of every single bus passenger every day.

Even 100 per cent checking will not always guarantee that all defects or problems will be identified:

- Checks may be inherently difficult. Although a doctor may undertake all the correct testing procedures to check for a particular disease, he or she may not necessarily be certain to diagnose it.
- Staff may become fatigued. For example, when inspecting repetitive items where it is easy to make mistakes. (Try counting the number of 'e's on this page. Count them again and see if you get the same score.)
- Information may be unreliable. Although all the customers in a restaurant may tell the head waiter that 'everything is fine', they may actually have reservations about their experience.

Sometimes, however, it is necessary to sample everything that is produced by a process or an operation. If a product is so critical that its failure to conform to specification would result in death or injury (for example, some parts that are used in aircraft, or some services within health care operations) then, although expensive, one hundred per cent inspection is necessary. In such cases, the consequence of non-conformance drives the decision to inspect everything. In other cases, it may be that the economics of a hundred per cent inspection are such that the cost of doing it is relatively small. For example, some labels can be automatically scanned as they are produced at virtually no extra cost. Yet, whenever a hundred per cent inspection is adopted, there is another risk – that of classifying something as an error when, in fact, it conforms to the specification. This distinction is summarised in what is often referred to as type I and type II errors.

Type I and type II errors

Checking quality by sampling, although requiring less time than checking everything, does have its own problems. Take the example of someone waiting to cross a street. There are two main options: cross (take some action), or continue waiting (take no action). If there is a break in the traffic and the person crosses, or if that person continues to wait because the traffic is too dense, then a correct decision has been made to cross (the action was appropriate for the circumstances). There are two types of incorrect decisions, or errors. One would be a decision to cross (take some action) when there is not an adequate break in the traffic, resulting in an accident; this is referred to as a type I error. Another would be a decision not to cross even though there was an adequate gap in the traffic; this is a type II error. Type I errors are those which occur when a decision was made to do something and the situation did not warrant it. Type II errors are those that occur when nothing was done, yet a decision to do something should have been taken as the situation did indeed warrant it. So, there are four outcomes, summarised in Table 13.4.

Table 13.4 **Type I and type II errors for a pedestrian crossing the road**

		Road conditions	
		Safe *(action was appropriate)*	**Unsafe** *(action was not appropriate)*
Decision	**Cross** (take some action)	Correct decision	**Type I error**
	Wait (take no action)	**Type II error**	Correct decision

Statistical process control

The most common method of checking the quality of a sampled product or service to make inferences about all the output from a process is called statistical process control (SPC). SPC is concerned with sampling the process during the production of the goods or the delivery of service. Based on this sample, decisions are made about whether the process is 'in control', that is, operating as it should be. If there seems to be a problem with the process, then it can be stopped (if possible and appropriate) and the problem identified and rectified. For example, an international airport may regularly ask a sample of customers if the cleanliness of its restaurants is satisfactory. If an unacceptable number of customers in one sample are found to be unhappy, airport managers may have to consider improving the procedures in place for cleaning tables.

Control charts

The value of SPC is not just to make checks of a single sample, but to monitor the results of many samples over a period of time. It does this by using control charts. Control charts record some aspect of quality (or performance generally) over time to see if the process seems to be performing as it should (called *in control*), or not (called *out of control*). If the process does seem to be going out of control, then steps can be taken *before* there is a problem.

Figure 13.6 shows a typical control chart. Charts something like these, can be found in almost any operation. They could, for example, represent the percentage of customers in a sample of 1,000 who, each week, were dissatisfied with the service they received from two call centres. In chart (a), measured customer dissatisfaction has been steadily increasing over time. There is evidence of a clear (negative) trend that management may wish to investigate. In chart (b), although there is little evidence of any trend in average dissatisfaction, the variability in performance seems to be increasing. Again the operation may want to investigate the causes.

Looking for *trends* is an important use of control charts. If the trend suggests the process is getting steadily worse, then it will be worth investigating the process. If the trend is steadily improving, it may still be worthy of investigation to try to identify what is happening that is making the process better. An even more important use of control charts is to investigate the *variation* in performance.

Why variation is a bad thing

Although a trend such as that shown in Figure 13.6(a) clearly indicates deteriorating performance, the variation shown in Figure 13.6(b) can be just as serious. Variation is a problem because it masks any changes in process behaviour. Figure 13.7 shows the performance of two

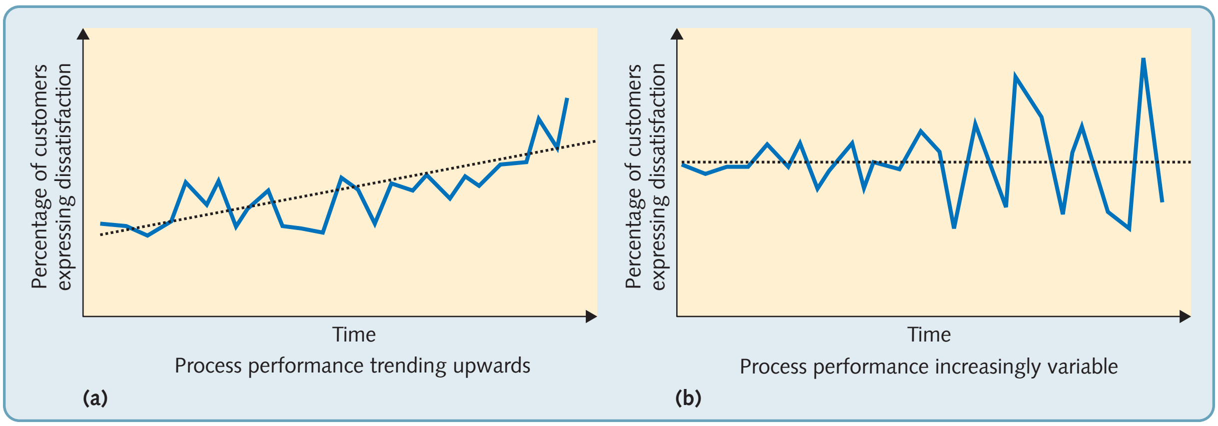

Figure 13.6 Control charting – any aspect of the performance of a process is measured over time and may show trends in average performance, and/or changes in the variation of performance over time

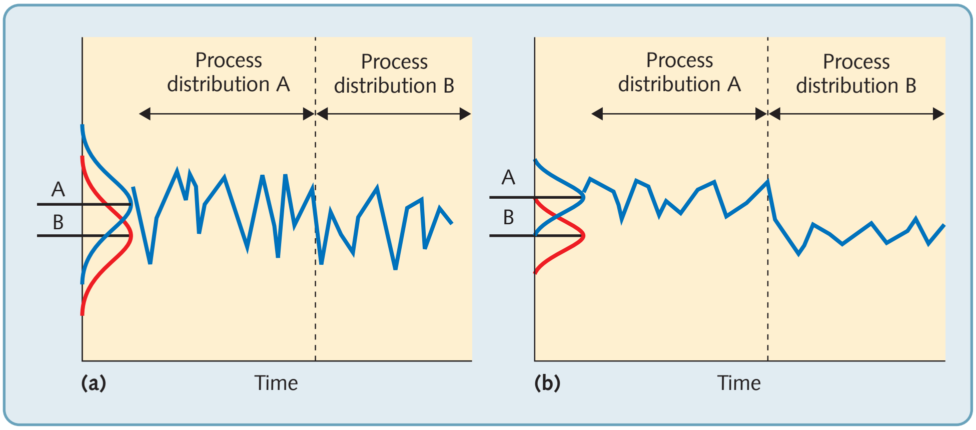

Figure 13.7 Low process variation allows changes in process performance to be readily detected

processes both of which change their behaviour at the same time. The process on the left has such a wide natural variation that it is not immediately apparent that any change has taken place. Eventually, it will become apparent, but it may take some time. By contrast, the performance of process represented by the chart on the right has a far narrower band of variation, so the same change in average performance is more easily noticed. The narrower the variation of a process, the more obvious are any changes that might occur, and the easier it is to make a decision to intervene. SPC is discussed much further in the supplement to this chapter. It is also one of the core ideas in the Six-Sigma improvement approach that was discussed in Chapter 12.

OPERATIONS PRINCIPLE

High levels of variation reduce the ability to detect changes in process performance.

EXAMPLE

What a giveaway[7]

Another negative effect of wide process variability is particularly evident in any process that fills any kind of container with weighed product. Even small reductions in the variability in filling levels can translate into major savings. This is because of what is known as 'giveaway' or 'over-fill', caused by the necessity to ensure that containers are not legally under-weight. Although slightly different regulations may apply in various parts of the world, any process that produces products that have an 'e' after the stated weight on the container must produce products with an average weight greater than the declared weight on the container, with the average weight being determined by sampling. In addition, there are two other legal conditions. First, no more than 2.5 per cent of the sample can lie between an upper and lower control limit. Second, no product under a lower control limit may be produced at all. Therefore, operations managers often build in a margin of safety in order to overcome that variation while allowing them to meet legal weights and measures conditions on minimum fill levels. As a result, containers on filling lines are routinely over-filled with more finished product than need be the case. This is why minimising fill variation can avoid 'giveaway'.

Process control, learning and knowledge

In recent years, the role of process control and SPC, in particular, has changed. Increasingly, it is seen not just as a convenient method of keeping processes in control, but also as an activity that is fundamental to the acquisition of competitive advantage. This is a remarkable shift in the status of SPC. Traditionally it was seen as one of the most *operational*, immediate and 'hands-on' operations management techniques. Yet, it is now seen as contributing to an operation's *strategic* capabilities. This is how the logic of the argument goes:

> OPERATIONS PRINCIPLE
>
> ***Statistical-based control gives the potential to enhance process knowledge.***

1. SPC is based on the idea that process variability indicates whether a process is in control or not.
2. Processes are brought into *control* and improved by progressively reducing process variability. This involves eliminating the assignable causes of variation.
3. One cannot eliminate assignable causes of variation without gaining a better understanding of how the process operates. This involves *learning* about the process, where its nature is revealed at an increasingly detailed level.
4. This learning means that *process knowledge* is enhanced, which in turn means that operations managers are able to predict how the process will perform under different circumstances. It also means that the process has a greater capability to carry out its tasks at a higher level of performance.
5. This increased *process capability* is particularly difficult for competitors to copy. It cannot be bought 'off-the-shelf'. It only comes from time and effort being invested in controlling operations processes. Therefore, process capability leads to strategic advantage.

In this way, process control leads to learning which enhances process knowledge and builds difficult-to-imitate process capability.

DIAGNOSTIC QUESTION

Does quality management always lead to improvement?

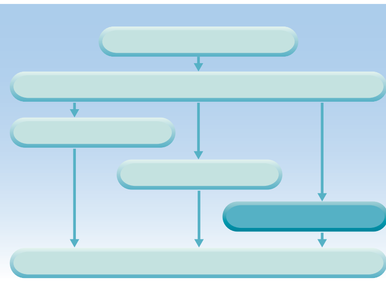

No amount of effort put into quality initiatives can guarantee improvement in process performance. In fact, some surveys show that up to half of quality programmes provide only disappointing, if any, permanent improvement. Improving quality is not something that happens simply by getting everyone in an organisation to 'think quality'. Very often improvements do not stick because there is no set of systems and procedures to support and embed them into the operation's day-to-day routines. 'Quality systems' are needed.

A quality system is '*the organizational structure, responsibilities, procedures, processes and resources for implementing quality management*'.[8] It should cover all facets of a business's operations and processes, and define the responsibilities, procedures and processes that ensure the implementation of quality improvement. The best-known quality system is ISO 9000.

The ISO 9000 approach

The ISO 9000 series is a set of worldwide standards that establishes requirements for companies' quality management systems. It is being used worldwide to provide a framework for quality assurance. By 2000, ISO 9000 had been adopted by more than a quarter of a million organisations in 143 countries. Originally, its purpose was to provide an assurance to the purchasers of products or services by defining the procedures, standards and characteristics of the control system that governed the process that produced them. In 2000, ISO 9000 was substantially revised. Rather than using different standards for different functions within a business, it took a 'process' approach and focused on the outputs from any operation's process, rather than the detailed procedures that had dominated the previous version. This process orientation requires operations to define and record core processes and sub-processes. ISO 9000 (2000) also stresses four other principles:

- Quality management should be customer focused and customer satisfaction measured using surveys and focus groups. Improvement against customer standards should be documented.
- Quality performance should be measured, and relate to products and services themselves, the processes that created them and customer satisfaction. Furthermore, measured data should always be analysed.
- Quality management should be improvement driven. Improvement must be demonstrated in both process performance and customer satisfaction.
- Top management must demonstrate their commitment to maintaining and continually improving management systems. This commitment should include communicating the importance of meeting customer and other requirements, establishing a quality policy and quality objectives, conducting management reviews to ensure the adherence to quality policies and ensuring the availability of the necessary resources to maintain quality systems.

The Deming Prize

The Deming Prize was instituted by the Union of Japanese Scientists and Engineers in 1951 and is awarded to those companies, initially in Japan, but more recently opened to overseas companies, which have successfully applied 'company-wide quality control' based upon statistical quality control. There are 10 major assessment categories: policy and objectives; organisation and its operation; education and its extension; assembling and disseminating of information; analysis; standardisation; control; quality assurance; effects; and future plans. The applicants are required to submit a detailed description of quality practices. This is a significant activity in itself and some companies claim a great deal of benefit from having done so.

The Malcolm Baldrige National Quality Award

In the early 1980s the American Productivity and Quality Center recommended that an annual prize, similar to the Deming Prize, should be awarded in America. The purpose of the awards was to stimulate American companies to improve quality and productivity, to recognise achievements, to establish criteria for a wider quality effort and to provide guidance on quality improvement. The main examination categories are: leadership, information and analysis, strategic quality planning, human resource utilisation, quality assurance of products and services, quality results and customer satisfaction. The process, like that of the Deming Prize, includes a detailed application and site visits.

The EFQM Excellence Model

Over 20 years ago Western European companies formed the European Foundation for Quality Management (EFQM). Since then the importance of quality excellence has become far more accepted. According to the EFQM: '*Whilst there are numerous management tools and techniques commonly used, the EFQM Excellence Model provides an holistic view of the organisation and* *it can be used to determine how these different methods fit together and complement each other. The model . . . .* [is] *. . . an overarching framework for developing sustainable excellence. Excellent organisations achieve and sustain outstanding levels of performance that meet or exceed the expectations of all their stakeholders. The EFQM Excellence Model allows people to understand the cause and effect relationships between what their organisation does and the results it achieves.'*[9]

The model is based on the idea that it is important to understand the cause and effect relationships between what an organisation does, (what it terms 'the Enablers'), and the results it achieves. The EFQM Excellence Model is shown in Figure 13.8. There are five enablers:

1. ***Leadership:*** that looks to the future, acts as a role model for values and ethics, inspires trust, is flexible, enables anticipation and so can react in a timely manner.
2. ***Strategy:*** that implements the organisation's mission and vision by developing and deploying a stakeholder-focused strategy.
3. ***People:*** organisations should value their people, creating a culture that allows mutually beneficial achievement of both organisational and personal goals, develops the capabilities of people, promotes fairness and equality, cares for, communicates, rewards and recognises people, in a way that motivates and builds commitment.
4. ***Partnership and resources:*** organisations should plan and manage external partnerships, suppliers and internal resources in order to support strategy and policies and the effective operation of processes.
5. ***Processes, products and services:*** organisations should design, manage and improve processes to ensure value for customers and other stakeholders.

Results are assessed using four criteria. They are:

1. ***Customer results:*** meeting or exceeding the needs and expectations of customers.
2. ***People results:*** meeting or exceeding the needs and expectations of employees.
3. ***Society results:*** achieving and sustaining results that meet or exceed the needs and expectations of the relevant stakeholders within society.
4. ***Business results:*** achieving and sustaining results that meet or exceed the needs and expectations of business stakeholders.

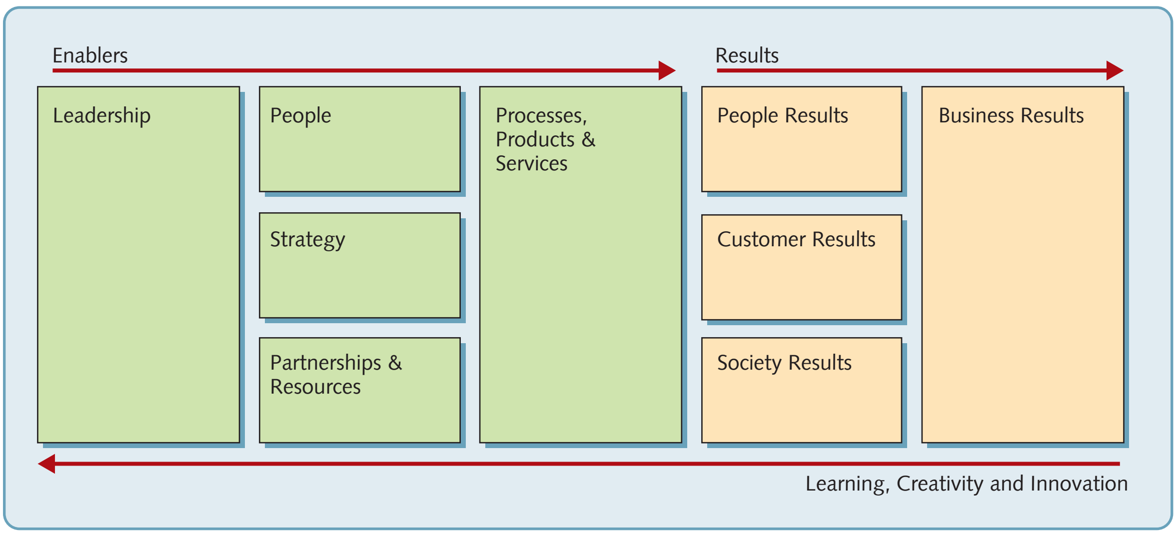

Figure 13.8 The EFQM Excellence Model (reproduced with the permission of the EFQM)

Critical commentary

Quality management has been one of the hottest topics in operations management and one of the most controversial. Much of the debate has centred on the people focus of quality management, especially the rhetoric of employee empowerment central to several modern approaches to quality. In many cases, it can be little more than an increase in employee discretion over minor details of their working practice. Some industrial relations academics argue that TQM rarely affects the fundamental imbalance between managerial control and employees' influence over organisational direction. For example, '... there is little evidence that employee influence over corporate decisions which affect them has been, or can ever be, enhanced through contemporary configuration of involvement. In other words, whilst involvement might increase individual task discretion, or open up channels for communication, the involvement programme is not designed to offer opportunities for employees to gain or consolidate control over the broader environment in which their work is located.'[10]

- Other criticisms concern the appropriateness of some mechanisms such as service level agreements (SLAs). Some see the strength of SLAs as the degree of formality they bring to customer–supplier relationships, but there also drawbacks. The first is that the 'pseudo-contractual' nature of the formal relationship can work against building partnerships. This is especially true if the SLA includes penalties for deviation from service standards. The effect can sometimes be to inhibit rather than encourage joint improvement. The second is that SLAs, tend to emphasise the 'hard' and measurable aspects of performance rather than the 'softer' but often more important aspects. A telephone might be answered within four rings, for example, but how the caller is treated, in terms of 'friendliness' may be far more important.

- Similarly, and notwithstanding its widespread adoption (and its revision to take into account some of its perceived failing), ISO 9000 is not seen as beneficial by all authorities. Criticisms include the following:

  - The whole process of documenting processes, writing procedures, training staff and conducting internal audits is expensive and time-consuming.
  - Similarly, the time and cost of achieving and maintaining ISO 9000 registration are excessive.
  - It is too formulaic. It encourages operations to 'manage by manual', substituting a 'recipe' for a more customised and creative approach to managing operations improvement.

SUMMARY CHECKLIST

- ☐ Does everyone in the business really believe in the importance of quality, or is it just one of those things that people say without really believing it?
- ☐ Is there an accepted definition of quality used within the business?
- ☐ Do people understand that there are many different definitions and approaches to quality, and do they understand why the business has chosen its own particular approach?
- ☐ Do all parts of the organisation understand their contribution to maintaining and improving quality?
- ☐ Are service level agreements used to establish concepts of internal customer service?
- ☐ Is some form of gap model used to diagnose quality problems?
- ☐ Is quality defined in terms of a series of quality characteristics?
- ☐ Is quality measured using all relevant quality characteristics?
- ☐ Is the cost of quality measured?
- ☐ Are quality costs categorised as prevention, appraisal, internal failure and external failure costs?
- ☐ Is quality adequately controlled?
- ☐ Has the idea of statistical process control (SPC) been explored as a mechanism for controlling quality?
- ☐ Do individual processes have any idea of their own variability of quality performance?
- ☐ Have quality systems been explored, such as ISO 9000 and the EFQM Excellence Model?

CASE STUDY

Turnround at the Preston Plant

'Before the crisis the quality department was just for looks, we certainly weren't used much for problem solving, the most we did was inspection. Data from the quality department was brought to the production meeting and they would all look at it, but no one was looking behind it.' (Quality Manager, Preston Plant)

The Preston plant of Rendall Graphics was located in Preston, Vancouver, across the continent from their headquarters in Massachusetts. Rendall bought the plant from the Georgetown Corporation in March 2000. Precision-coated papers for ink-jet printers accounted for the majority of the plant's output, especially paper for specialist uses. The plant used coating machines that allowed precise coatings to be applied. After coating, the conversion department cut the coated rolls to the final size and packed the sheets in small cartons.

The curl problem

In late 1998 Hewlett Packard (HP), the plant's main customer for ink-jet paper, informed the plant of some problems it had encountered with paper curling under conditions of low humidity. There had been no customer complaints to HP, but their own personnel had noticed the problem, and they wanted it fixed. Over the next seven or eight months a team at the plant tried to solve the problem. Finally, in October of 1999 the team made recommendations for a revised and considerably improved coating formulation. By January 2000, the process was producing acceptably. However, 1999 had not been a good year for the plant. Although sales were reasonably buoyant, the plant was making a loss of around $2 million for the year. In October 1999, Tom Branton, previously accountant for the business, was appointed as managing director.

Slipping out of control

In the spring of 2000, productivity, scrap and rework levels continued to be poor. In response to this, the operations management team increased the speed of the line and made a number of changes to operating practice in order to raise productivity.

'Looking back, changes were made without any proper discipline, and there was no real concept of control. We were always meeting specification, yet we didn't fully understand how close we really were to not being able to make it. The culture here said, "If it's within specification then it's OK" and we were very diligent in making sure that the product which was shipped ***was*** *in specification. However, Hewlett Packard gets "process charts" that enables them to see more or less exactly what is happening right inside your operation. We were also getting all the reports but none of them were being internalised; we were using them just to satisfy the customer. By contrast, HP have a statistically based analytical mentality that says to itself, "You might be capable of making this product but we are thinking two or three product generations forward and asking ourselves, will you have the capability then, and do we want to invest in this relationship for the future?"* '(Tom Branton)

The spring of 2000 also saw two significant events. First, Hewlett Packard asked the plant to bid for the contract to supply a new ink-jet platform, known as the Vector project, a contract that would secure healthy orders for several years. The second event was that the plant was acquired by Rendall.

"What did Rendall see when they bought us? They saw a small plant on the Pacific coast losing lots of money." (Finance Manager, Preston Plant)

Rendall were not impressed by what they found at the Preston plant. It was making a loss and had only just escaped from incurring a major customer's disapproval over the curl issue. If the plant did not get the Vector contract, its future looked bleak. Meanwhile the chief concern continued to be productivity. Once again, there were occasional complaints about quality levels. However, HP's attitude caused some bewilderment to the operations management team.

'When HP asked questions about our process the operations guys would say, "Look we're making roll after roll of paper, it's within specification. What's the problem?"' (Quality Manager, Preston Plant)

It was not until summer that the full extent of HP's disquiet was made. *'I will never forget June of 2000. I was at a meeting with HP in Chicago. It was not even about quality. But during the meeting one of their engineers handed me a control chart, one that we supplied with every batch of product. He said "Here's your latest control chart. We think you're out of control and you don't know that you're out of control and we think that we are looking at this data more than you are." He was absolutely right, and I fully understood how serious the position was. We had our most important customer telling us we couldn't run our processes just at the time we were trying to persuade them to give us the Vector contract.'* (Tom Branton)

The crisis

Tom immediately set about the task of bringing the plant back under control. They first decided to go back to the conditions that prevailed in January, when the curl team's recommendations had been implemented. This was the state before productivity pressures had caused the process to be adjusted. At the same time the team worked on ways of implementing unambiguous 'shut-down rules' that would allow operators to decide under what conditions a line should be halted if they were in doubt about the quality of the product they were making.

'At one point in May of 2000 we had to throw away 64 jumbo rolls of out-of-specification product. That's over $100,000 of product scrapped in one run. Basically, that was because they had been afraid to shut the line down. Either that or they had tried to tweak the line while it was running to get rid of the defect. The shut-down guidelines in effect say, "We are not going to operate when we are not in a state of control". Until then our operators just couldn't win. If they failed to keep the machines running we would say, "You've got to keep productivity up". If they kept the machines running but had quality problems as a result, we criticised them for making garbage. Now you get into far more trouble for violating process procedures than you do for not meeting productivity targets.' (Engineer, Preston Plant)

This new approach needed to be matched by changes in the way the communications were managed in the plant.

'We did two things that we had never done before. First each production team started holding daily reviews of control chart data. Second, one day a month we took people away from production and debated the control chart data. Several people got nervous because we were not producing anything. But it was necessary. For the first time you got operators from the three shifts meeting together and talking about the control chart data and other quality issues. Just as significantly, we invited HP up to attend these meetings. Remember these weren't staged meetings, it was the first time these guys had met together and there was plenty of heated discussion, all of which the Hewlett Packard representatives witnessed.' (Engineer, Preston Plant)

At last, something positive was happening in the plant and morale on the shop floor was buoyant. By September 2000, the results of the plant's teams efforts were starting to show results. Processes were coming under control, quality levels were improving and, most importantly, personnel both on the shop floor and in the management team were beginning to get into the 'quality mode' of thinking. Paradoxically, in spite of stopping the line periodically, the efficiency of the plant was also improving.

Yet the Preston team did not have time to enjoy their emerging success. In September of 2000, the plant learned that it would not get the Vector project because of their recent quality problems. Then Rendall decided to close the plant. *'We were losing millions, we had lost the Vector project, and it was really no surprise. I told the senior management team and said that we would announce it probably in April of 2001. The real irony was that we knew that we had actually already turned the corner.'* (Tom Branton)

Notwithstanding the closure decision, the management team in Preston set about the task of convincing Rendall, that the plant could be viable. They figured it would take three things. First, it was vital that they continue to improve quality. Progressing with their quality initiative involved establishing full statistical process control (SPC).

Second, costs had to be brought down. Working on cost reduction was inevitably going to be painful. The first task was to get an understanding of what should be an appropriate level of operating costs. *'We went through a zero-based assessment to decide what an ideal plant would look like, and the minimum number of people needed to run it.'* (Tom Branton)

By December of 2000, there were 40 per cent fewer people in the plant than two months earlier. All departments were affected. The quality department shrank more than most, from 22 people down to 6. *'When the plant was considering down-sizing they asked me, "How can we run a lab with six technicians?" I said, "Easy. We just make good paper in the first place, and then we don't have to* *inspect all the garbage. That alone would save an immense amount of time."'* (Quality Manager, Preston Plant)

Third, the plant had to create a portfolio of new product ideas, which could establish a greater confidence in future sales. Several new ideas were under active investigation, the most important of which was 'Protowrap', a wrap for newsprint that could be re-pulped. Although it was a technically difficult product, the plant's newly acquired capabilities allowed the product to be made economically.

Out of the crisis

In spite of their trauma, the plant's management team faced Christmas of 2000 with increasing optimism. They had just made a profit for the first time for over two years. By spring of 2001 even HP, at a corporate level, were starting to take notice. It was becoming obvious that the Preston plant really had made a major change. More significantly, HP had asked the plant to bid for a new product. April 2001 was a good month for the plant. It had chalked up three months of profitability and HP formally gave the new contract to Preston. Also in April, Rendall reversed their decision to close the plant.

QUESTIONS

1 What are the most significant events in the story of how the plant survived because of its adoption of quality-based principles?

2 The plant's processes eventually were brought under control. What were the main benefits of this?

3 SPC is an operational level technique of ensuring quality conformance. How many of the benefits of bringing the plant under control would you class as strategic?

APPLYING THE PRINCIPLES

Some of these exercises can be answered by reading the chapter. Others will require some general knowledge of business activity and some might require an element of investigation. **Model answers for the first two questions are available on the student companion website.** For model answers to the other questions in this section, please ask your tutor/lecturer.

1 Human error is a significant source of quality problems. Think through the times that (with hindsight) you have made an error and answer the following questions.

(a) How do you think that human error causes quality problems?

(b) What could one do to minimise human error?

2 The owner of a small wedding photography business realises that the market is changing. 'I used to take a few photos during the wedding ceremony and then formal group shots outside. It rarely took more than two hours. Around 30 photos would go in a standard wedding album. You had to get the photos right; that was really the only thing I was judged on. Now it's different. I spend all day at a wedding, and sometimes late into the evening. You're almost like another guest. All the guests at the wedding now are important. You have to get the best photos while being as discreet as possible. Clients judge you on both the pictures and the way you interact with everyone on the day. The product has changed too. Clients receive a memory stick with around 500 photos and a choice of 10 albums. I also offer photo books with greater customisation. I can offer albums with items such as invitations, confetti, and menus; and individual paintings created from photographs. Obviously, I would have to outsource the paintings. Wedding guests can order photos and related products online. My anxiety is that advertising this service at the wedding will be seen as being too commercial. We have a high level of demand in summer, with weekends booked up two years in advance. I may take on

additional photographers during busy periods, but the best ones are busy themselves. Also the business is about client relations and that's hard to replicate. I often offer clients advice on such things as locations, bands, caterers, and florists. Wedding planning is clearly an area that could be profitable to the business. Another option is to move beyond weddings into other areas, such as school photos, birthdays, celebrations, or studio work.'

(a) How has the business changed over time?

(b) What do you think are the key quality challenges facing the business?

(c) What do you think should be done to ensure the business maintains quality levels in the future?

3 Ryanair, the best-known budget airline in Europe, focuses on popular routes and very low operating costs. For years, the boss of the airline, Michael O'Leary's policy on customer service was clear. 'Our customer service', he said 'is about the most well defined in the world. We guarantee to give you the lowest airfare. You get a safe flight. You get a normally on-time flight. That's the package. We don't and won't give you anything more. Are we going to say sorry for our lack of customer service? Absolutely not. If a plane is cancelled, will we put you up in a hotel overnight? Absolutely not. If a plane is delayed, will we give you a voucher for a restaurant? Absolutely not.'

However, the bad publicity eventually prompted a limited rethink by the company. After a drop in their hitherto rapid profit growth, shareholder concern and Ryanair being voted the worst of Britain's 100 biggest brands by readers of the consumer magazine *Which?* The company announced that it was to reform its 'abrupt culture, and try to eliminate things that unnecessarily annoy customers'. Included in these annoying practices were fines for small luggage size transgressions and an unpopular €70 fee for issuing boarding passes at the airport, rather than printing it out at home (it was lowered to €10). Yet Ryanair insisted that such charges were not money-spinning schemes, but were designed to encourage operational efficiency that kept fairs low. In fact, fewer than ten passengers a day had to pay for forgotten boarding passes.

What does this example tell us about the trade-off between service quality and cost?

4 Understanding type I and type II errors is essential for a surgeon's quality planning. Consider, for example, appendectomy operations. Following a diagnosis of appendicitis, removal of the appendix is necessary because of the risk of it bursting, causing potentially fatal poisoning of the blood. The surgical procedure is a relatively simple operation, but there is always a small risk with any invasive surgery. It is also an expensive procedure; in the USA around $4,500 per operation. Unfortunately, appendicitis is difficult to diagnose and the diagnosis is only 10 per cent accurate. However, a new technique claims to be able to identify 100 per cent of true appendicitis cases prior to surgery; this costs less than $250, which means that one single avoided surgery pays for around 20 tests.

(a) How does this new test change the likelihood of type I and type II errors?

(b) Why is this important?

5 Look again at the example, 'Tea and Sympathy'.

(a) Why do you think 'Nicky's Rules' help to make the Tea and Sympathy operation more efficient?

(b) The restaurant's approach to quality of service seems very different to most restaurants. Why do you think it seems to work here?

6 Look again at the example, 'The Four Seasons Hotel Canary Wharf'.

(a) The company has what it calls its Golden Rule; 'Do to others (guests and staff) as you would wish others to do to you'. Why is this important in ensuring high-quality service?

(b) What do you think the hotel's guests expect from their stay?

(c) How do staff who use their own initiative contribute to quality?

Notes on chapter

1 Sources include: Vitaliev, V. (2009) 'The much-loved knife', *Engineering and Technology Magazine,* 21 July; Victorinox website (2012) 'The Victorinox Quality System'.

2 Interview with Michael Purtill, General Manager, Four Seasons Canary Wharf Hotel.

3 Sources include: *The Economist* (2013) 'Overtired and overdrawn', Economist Print Edition, 15 June; Wilson, H. (2014) 'Fat-fingered trader sets Tokyo alarms ringing', *The Times,* 2 October.

4 Zeithaml, V.A., Parasuraman, A. and Berry, L. (1990), Delivering Quality Service; Balancing Customer Perceptions and Expectations, *Free Press.*

5 Mechling, L. (2002) 'Get Ready for a storm in a tea shop', *The Independent,* 8 March and company web site (2011).

6 Source: Plunkett, J.J. and Dale, B.S. (1987) 'A review of the literature in quality-related costs', *The International Journal of Quality and Reliability Management,* vol. 4, no. 1.

7 Wheatley, M. (2010) 'Filling time on the production line', *Engineering and Technology* magazine, 8 November.

8 Dale, B.G, van der Wiele, T. and van Iwaarden, J. (2007) *Managing Quality,* (5th edn), Wiley-Blackwell.

9 The EFQM Website, www.efqm.org

10 Hyman, J. and Mason, B. (1995) *Management Employees Involvement and Participation,* Sage.

TAKING IT FURTHER

Crosby, P.B. (1979) *Quality is Free, McGraw-Hill.* *One of the gurus. It had a huge impact in its day. Read it if you want to know what all the fuss was about.*

Dale, B. G., van der Wiele, T. and van Iwaarden, J. (2007) *Managing Quality (5th edn), Wiley-Blackwell.* *This is the fifth edition of a book that has long been one of the best-respected texts in the area. A comprehensive and balanced guide to the area.*

Goetsch, D.L. and Davis, S. (2013) *Quality Management for Organizational Excellence: Introduction to Total Quality (7th edn), Pearson Education.* *Up-to-date account of the topic.*

Feigenbaum, A.V. (1986) *Total Quality Control, McGraw-Hill.* *A more comprehensive book than those by some of the other quality gurus.*

Hoyle, D. (2006) *Quality Management Essentials, Butterworth-Heinemann, Practical.*

Oakland, J. S. (2014) *Total Quality Management and Operational Excellence: Text with Cases, Pearson.* *The latest text from one of the gurus of quality.*

Pande, P.S., Neuman, R.P. and Kavanagh, R.R. (2000) *The Six Sigma Way, McGrawHill.* *There are many books written by consultants for practicing managers on the now fashionable Six Sigma Approach (see supplement to chapter). This one is readable and informative.*

SUPPLEMENT TO CHAPTER 13

Statistical process control (SPC)

Introduction

The purpose of statistical process control (SPC) is both to control process performance, keeping it within acceptable limits, and to improve process performance by reducing the variation in performance from its target level. It does this by applying statistical techniques to understand the nature of performance variation over time. For those who are anxious about the 'statistical' part of SPC, don't be. Essentially, SPC is based on principles that are both practical and intuitive. The statistical element is there to help rather than complicate quality decisions.

Variation in process performance

The core instrument in SPC is the control chart. (An explanation of this is given earlier in the chapter.) They are an illustration of the dynamic performance of a process, measuring how some aspect of process performance varies over time. All processes vary to some extent. No machine will give precisely the same result each time it is used. All materials vary a little. People in the process differ marginally in how they do things each time they perform a task. Given this, it is not surprising that any measure of performance quality (whether attribute or variable) will also vary. Variations that derive from these *normal* or *common causes* of variation can never be entirely eliminated (although they can be reduced).

For example, at a call centre for a utility company, customer care operatives answer queries over accounts, service visits, and so on. The length of time of each call will vary depending on the nature of the enquiry and the customer's needs. There will be some variation around an average call time. When the process of answering and responding to customer enquiries is stable, the computer system that intercepts and allocates calls to customer care operatives could be asked to randomly sample the length of each call. As these data built up, the histogram showing call times could develop as is shown in Figure 13.9. The first calls could lie anywhere within the natural variation of the process but are more likely to be close to the average call length (Figure 13.9(a)). As more calls are measured, they would clearly show a tendency to be close to the process average (see Figure 13.9(b) and (c)). Eventually, the data will show a smooth histogram that can be drawn into a smoother distribution that will indicate the underlying process variation (the distribution shown in Figure 13.9(f)).

Often this type of variation can be described by a normal distribution. (Even if these raw data do not conform to a normal distribution, they can be manipulated to approximate to one by using sampling, see later.) It is a characteristic of normal distributions that 99.7 per cent of the measures will lie within ± 3 standard deviations of the distribution (standard deviation is a measure of how widely the distribution is spread or *dispersed*).

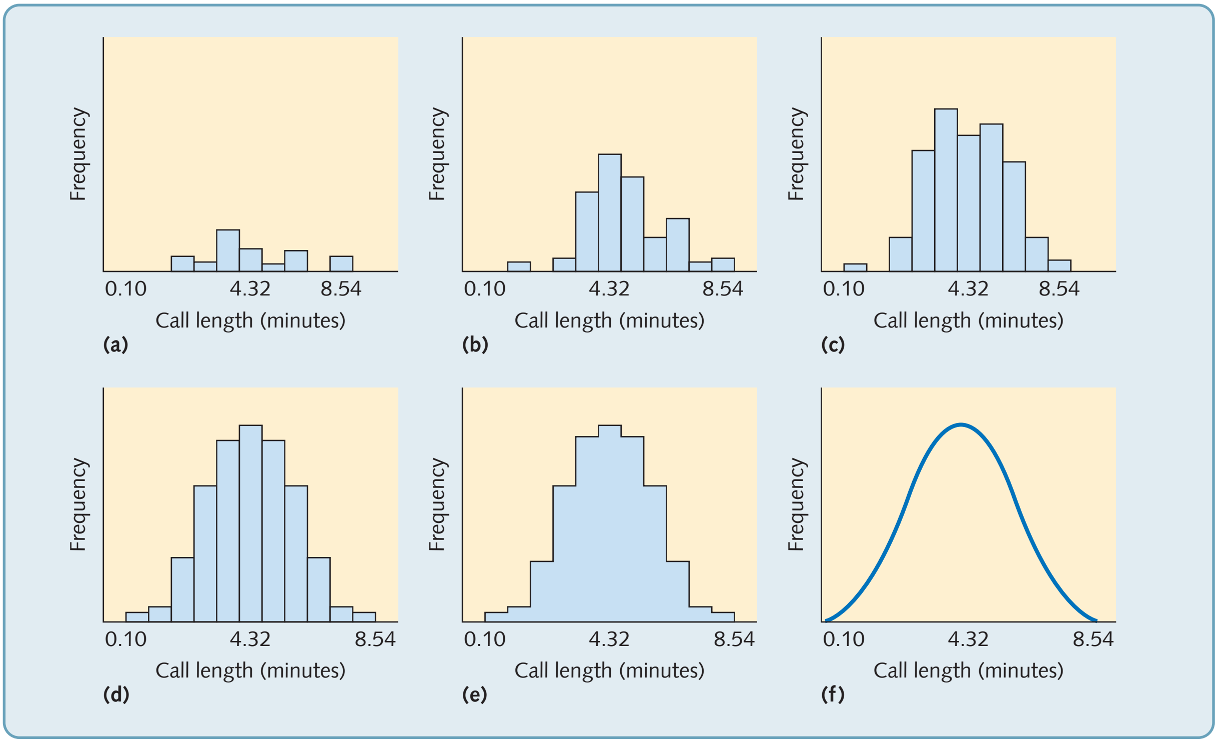

Figure 13.9 **The natural variation of call times in a call centre can be described by a normal distribution**

The central limit theorem

Not all processes will vary in their performance according to a normal distribution. However, if a sample is taken from any type of distribution, the distribution of the average of the sample (sample mean) **will** approximate to a normal distribution. For example, there is an equal probability of any number between one and six being thrown on a six-sided, unweighted die. The distribution is rectangular with an average of 3.5 as shown in Figure 13.10(a). But if a die is thrown (say) six times repeatedly and the average of the six throws calculated, the sample average will also be 3.5, but the standard deviation of the distribution will be the standard deviation of the original rectangular distribution divided by the square of the sample size. More significantly, the shape of the distribution will be close to normal and so can be treated the same way as a normal distribution. This becomes important when control limits are calculated (see later).

Is the process 'in control'?

Not all variation in process performance is the result of common causes. There may be something wrong with the process that is assignable to an abnormal and preventable cause. Machinery may have worn or been set up badly. An untrained person may not be following the prescribed procedure for the process. The causes of such variation are called *assignable or abnormal causes.* The question for operations management is whether the results from any particular sample, when plotted on the control chart, simply represent the variation due to

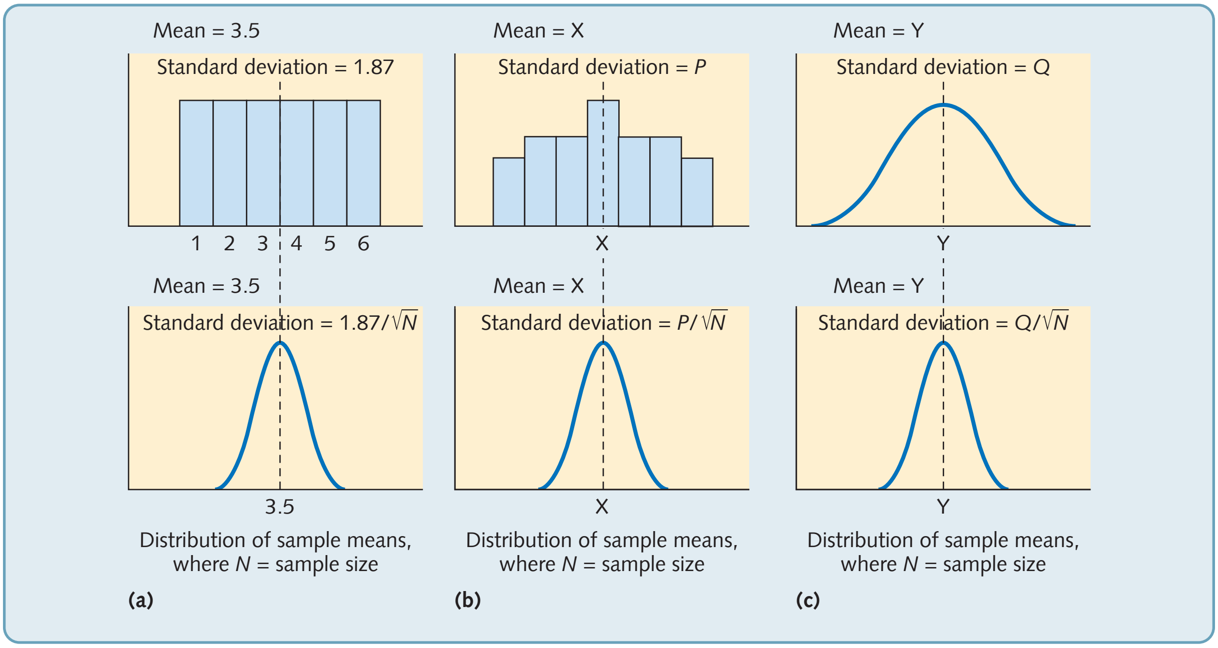

Figure 13.10 The distribution of sample means (averages) from any distribution will approximate to a normal distribution

common causes or due to some specific and correctable, *assignable* cause. Figure 13.11(a), for example, shows the control chart for the average call length of samples of customer calls in a utility's call centre. Like any process, the results vary, but the last three points seem to be lower than usual. The question is whether this is natural variation or the symptom of some more serious cause. Is the variation the result of common causes, or does it indicate assignable causes (something abnormal) occurring in the process?

To help make this decision, control limits can be added to the control charts that indicate the expected extent of 'common-cause' variation. If any points lie outside these control limits then the process can be deemed *out of control* in the sense that variation is likely to be due to assignable causes. These can be set in a statistically revealing manner based on the probability that the mean of a particular sample will differ by more than a set amount from the mean of the population from which it is taken. Figure 13.11(b) shows the same control chart as Figure 13.11(a) with the addition of control limits put at ±3 standard deviations (of the population of sample means) away from the mean of sample averages. It shows that the probability of the final point on the chart being influenced by an assignable cause is very high indeed. When the process is exhibiting behaviour that is outside its normal 'common-cause' range, it is said to be 'out of control'.

However, we cannot be absolutely certain that the process is out of control. There is a small but finite chance that the point is a rare but natural result at the tail of its distribution. Stopping the process under these circumstances would represent a type I error because the process is actually in control. Alternatively, ignoring a result which in reality is due to an assignable cause is a type II error (see Table 13.5). Control limits that are set at three standard deviations either side of the population mean are called the upper control limit (UCL) and lower control limit (LCL). There is only a 0.3 per cent chance of any sample mean falling outside these limits by chance causes (that is, a chance of a type I error of 0.3 per cent).

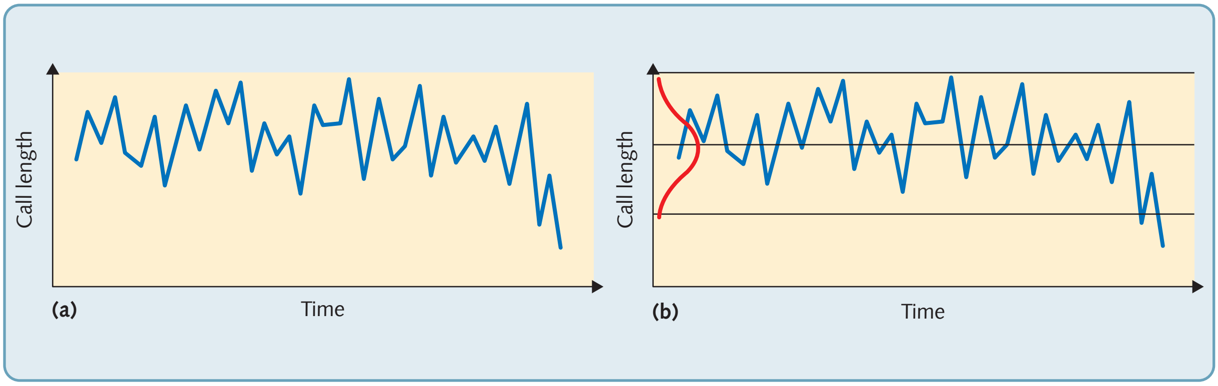

Figure 13.11 Control chart for the average call length in a call centre

Table 13.5 Type I and type II errors in SPC

		Actual process state	
		In control	*Out of control*
Decision	**Stop process**	**Type I error**	Correct decision
	Leave alone	Correct decision	**Type II error**

Process capability

Using control charts to assess whether the process is in control is an important internal benefit of SPC. An equally important question for any operations manager would be: 'Is the variation in the process performance acceptable to external customers?' The answer will depend on the acceptable range of performance that will be tolerated by the customers. This range is called the *specification range*. Returning to the boxed rice example, if the weight of rice in the box is too small then the organisation might infringe labeling regulations; if it is too large, the organisation is 'giving away' too much of its product for free.

Process capability is a measure of the acceptability of the variation of the process. The simplest measure of capability (C_p) is given by the ratio of the specification range to the 'natural' variation of the process (i.e. ± 3 standard deviations):

$$C_p = \frac{\text{UTL} - \text{LTL}}{6s}$$

where UTL = the upper tolerance limit
LTL = the lower tolerance limit
s = the standard deviation of the process variability.

Generally, if the C_p of a process is greater than 1, it is taken to indicate that the process is 'capable', and a C_p of less than 1 indicates that the process is not 'capable', assuming that the distribution is normal (see Figure 13.12(a), (b) and (c)).

The simple C_p measure assumes that the average of the process variation is at the mid-point of the specification range. Often the process average is offset from the specification range,

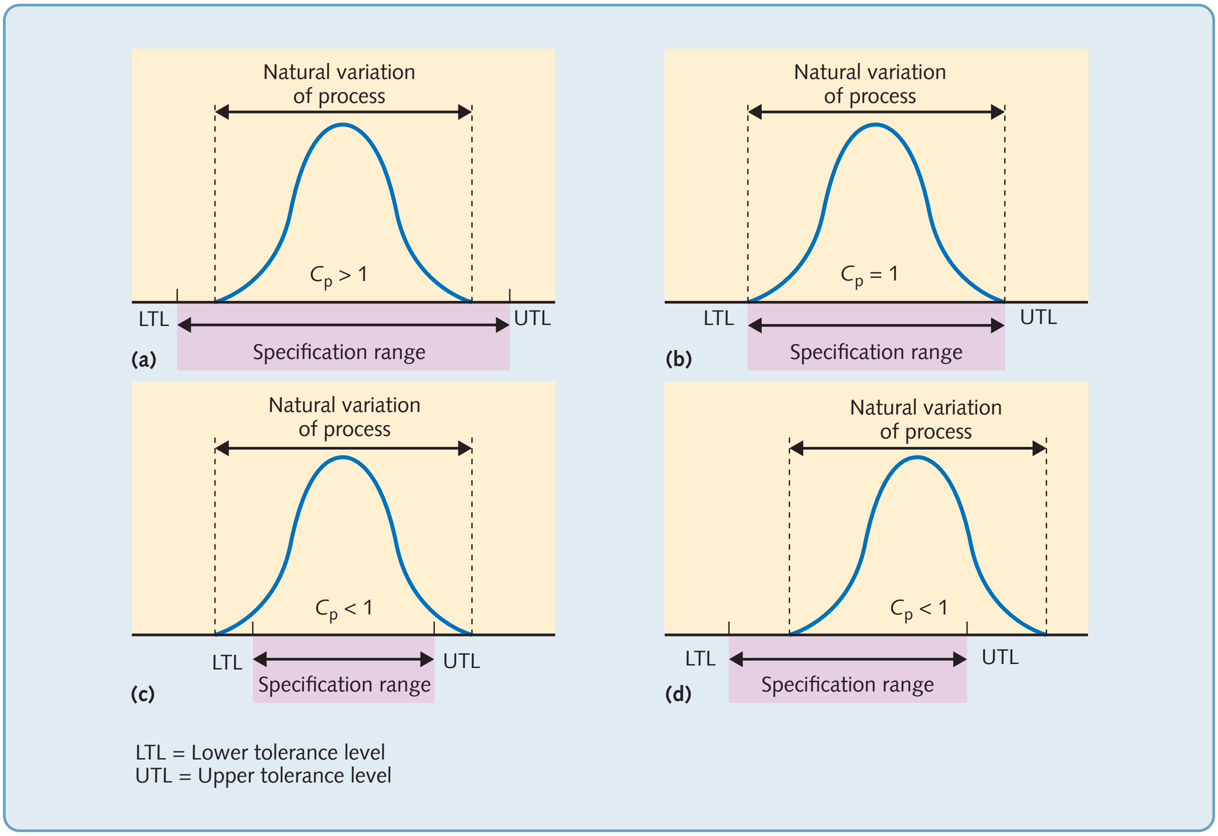

Figure 13.12 Process capability compares the natural variation of the process with the specification range that is required

however (see Figure 13.14(d)). In such cases, *one-sided* capability indices are required to understand the capability of the process:

$$\text{Upper one-sided index } C_{pu} = \frac{\text{UTL} - X}{3s}$$

$$\text{Lower one-sided index } C_{pl} = \frac{X - \text{LTL}}{3s}$$

Where X = the process average.

Sometimes only the lower of the two one-sided indices for a process is used to indicate its capability (C_{pk}):

$$C_{pk} = \min(C_{pu}, C_{pl})$$

WORKED EXAMPLE

In the case of the process filling boxes of rice, described previously, process capability can be calculated as follows:

$$\text{Specification range} = 214 - 198 = 16 \text{ g}$$

$$\text{Natural variation of process} = 6 \times \text{standard deviation}$$

$$= 6 \times 2 = 12 \text{ g}$$

$$C_p = \text{process capability}$$

$$= \frac{UTL - LTL}{6s}$$

$$= \frac{214 - 198}{6 \times 2} = \frac{16}{12}$$

$$= 1.333$$

If the natural variation of the filling process changed to have a process average of 210 grams but the standard deviation of the process remained at 2 grams:

$$C_{pu} = \frac{214 - 210}{3 \times 2} = \frac{4}{6} = 0.666$$

$$C_{pl} = \frac{210 - 198}{3 \times 2} = \frac{12}{6} = 2.0$$

$$C_{pk} = \min(0.666, 2.0)$$

$$= 0.666$$

Control charts for variables

The most commonly used type of control chart employed to control variables is the $\bar{X} - R$ *chart*. In fact, this is really two charts in one. One chart is used to control the sample average or mean ($\bar{X}$). The other is used to control the variation within the sample by measuring the range (R). The range is used because it is simpler to calculate than the standard deviation of the sample.

The means ($\bar{X}$) chart can pick up changes in the average output from the process being charted. Changes in the means chart would suggest that the process is drifting generally away from its supposed process average, although the variability inherent in the process may not have changed. The range (R) chart plots the range of each sample, which is the difference between the largest and the smallest measurement in the samples. Monitoring sample range gives an indication of whether the variability of the process is changing, even when the process average remains constant.

Control limits for variables control chart

As with attributes control charts, a statistical description of how the process operates under normal conditions (when there are no assignable causes) can be used to calculate control limits. The first task in calculating the control limits is to estimate the grand average or population mean ($\bar{\bar{X}}$) and average range ($\bar{R}$) using m samples each of sample size n.

The population mean is estimated from the average of a large number (m) of sample means:

$$\bar{\bar{X}} = \frac{\bar{X}_1 + \bar{X}_2 + \ldots \bar{X}_m}{m}$$

The average range is estimated from the ranges of the large number of samples:

$$\bar{R} = \frac{R_1 + R_2 + \ldots R_m}{m}$$

The control limits for the sample means chart are:

$$\text{Upper control limit (UCL)} = \bar{\bar{X}} + A_2\bar{R}$$

$$\text{Lower control limit (LCL)} = \bar{\bar{X}} - A_2\bar{R}$$

Table 13.6 Factors for the calculation of control limits

Sample size n	A_2	D_3	D_4
2	1.880	0	3.267
3	1.023	0	2.575
4	0.729	0	2.282
5	0.577	0	2.115
6	0.483	0	2.004
7	0.419	0.076	1.924
8	0.373	0.136	1.864
9	0.337	0.184	1.816
10	0.308	0.223	1.777
12	0.266	0.284	1.716
14	0.235	0.329	1.671
16	0.212	0.364	1.636
18	0.194	0.392	1.608
20	0.180	0.414	1.586
22	0.167	0.434	1.566
24	0.157	0.452	1.548

The control limits for the range charts are:

$$\text{Upper control limit (UCL)} = D_4\overline{R}$$

$$\text{Lower control limit (LCL)} = D_3\overline{R}$$

The factors A_2, D_3 and D_4 vary with sample size and are shown in Table 13.6.

The LCL for the means chart may be negative (for example, temperature or profit may be less than zero) but it may not be negative for a range chart (or the smallest measurement in the sample would be larger than the largest). If the calculation indicates a negative LCL for a range chart then the LCL should be set to zero.

WORKED EXAMPLE

GAM (Groupe As Maquillage) is a contract cosmetics company that manufactures and packs cosmetics and perfumes for other companies. One of its plants operates a filling line that automatically fills plastic bottles with skin cream and seals the bottles with a screw-top cap. The tightness with which the screw-top cap is fixed is an important aspect of quality. If the cap is screwed on too tightly, there is a danger that it will crack; if screwed on too loosely it might come loose. Either outcome could cause leakage. The plant had received some complaints of product leakage, possibly caused by inconsistent fixing of the screw-tops. Tightness can be measured by the amount of turning force (torque) that is required to unfasten the tops. The company decided to take samples of the bottles coming out of the filling-line process, test them for their unfastening torque and plot the results on a control chart. Several

GAM *Quality Counts*

PART NO. 73/70328A			CHART
PROCESS Torque Cap #3	PRODUCT Ski lan 707	TOLERANCE 807 – 817	
MACHINE M732	OPERATOR CH	UNITS g/cm³	

Date		2/6									
Time		6:00	6:30	7:00	7:30	8:00	8:30	9:00	9:30	10:00	10:30
Sample measurements	1	811	813	811	817	812	813	813	817	808	805
	2	812	814	813	814	813	814	815	813	813	808
	3	813	815	815	812	812	814	813	814	811	806
	4	812	814	813	813	811	815	811	815	813	807
	5										
Sum		3248	3256	3252	3248	3248	3256	3252	3259	3245	3225
Mean, X		812	814	813	814	812	814	813	815	811	807
Range, R		2	2	4	5	2	2	4	4	5	3

Date														
Time		11:00	11:30	12:00	12:30	13:00	13:30	14:00	14:30	15:00	15:30			
Sample measurements	1	813	812	811	808	812	808	808	804	804	813			
	2	812	805	813	806	813	809	818	813	812	819			
	3	816	806	811	803	808	813	819	812	812	813			
	4	809	808	809	805	804	814	821	808	820	807			
	5													
Sum		3250	3231	3244	3222	3237	3244	3266	3237	3248	3252			
Mean, X		813	808	811	806	809	811	817	809	812	813			
Range, R		7	7	4	5	9	6	13	9	16	12			

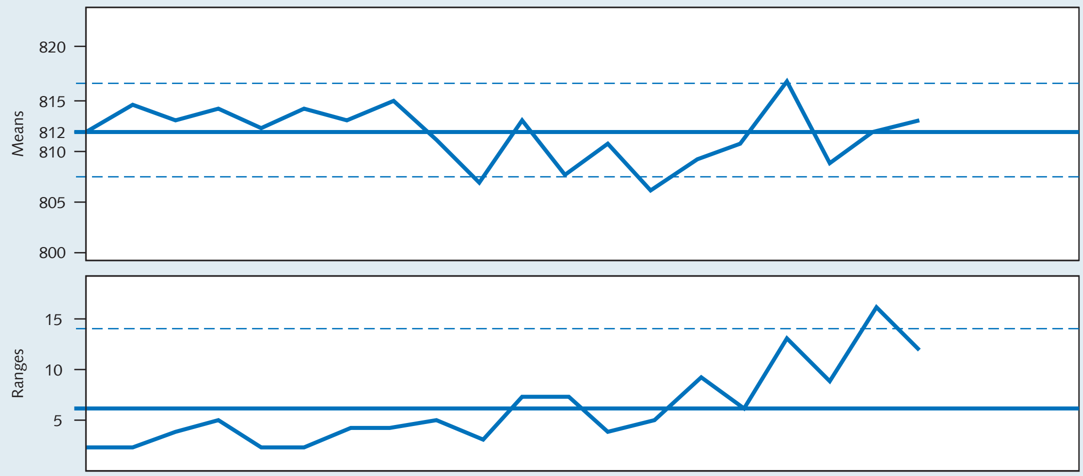

Figure 13.13 **The completed control form for GAM's torque machine showing the mean ($\overline{X}$) and range ($\overline{R}$) charts**

samples of four bottles are taken during a period when the process is regarded as being in control (Figure 13.13).

The following data are calculated from this exercise:

The grand average of all samples $\overline{\overline{X}} = 812 \text{ g/cm}^3$
The average range of the sample $\overline{R} = 6 \text{ g/cm}^3$

Control limits for the means ($\overline{X}$) chart were calculated as follows:

$$\text{UCL} = \overline{\overline{X}} + A_2\overline{R}$$
$$= 812 + (A_2 \times 6)$$

From Table 13.6, we know, for a sample size of four, $A_2 = 0.729$. Thus:

$$\text{UCL} = 812 + (0.729 \times 6)$$
$$= 816.37$$
$$\text{LCL} = \overline{\overline{X}} - (A_2\overline{R})$$
$$= 812 - (0.729 \times 6)$$
$$= 807.63$$

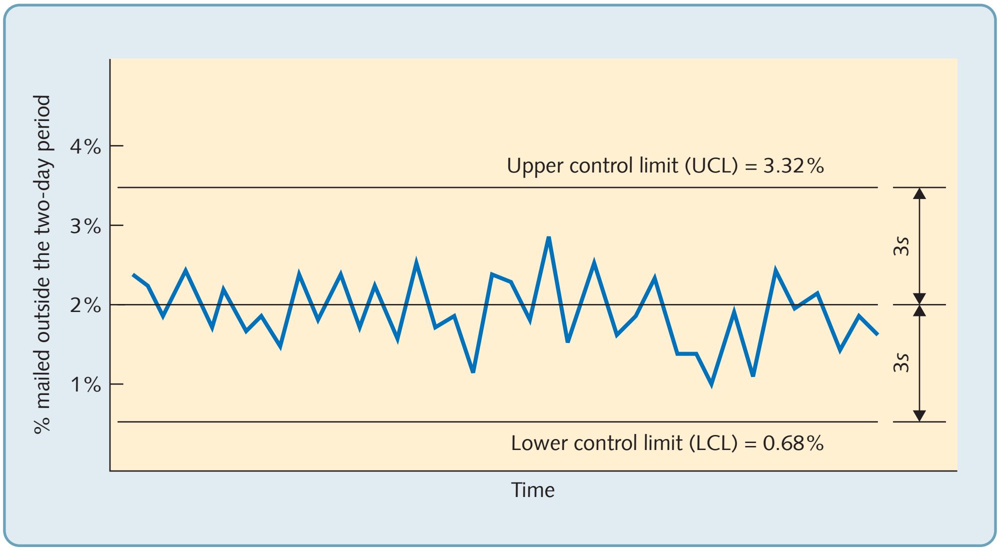

Figure 13.14 Control chart for the percentage of customer accounts which are mailed outside their two-day period

Control limits for the range chart (R) were calculated as follows:

$$\begin{aligned} \text{UCL} &= D_4 \times \overline{R} \\ &= 2.282 \times 6 \\ &= 13.69 \\ \text{LCL} &= D_3\overline{R} \\ &= 0 \times 6 \\ &= 0 \end{aligned}$$

After calculating these averages and limits for the control chart, the company regularly took samples of four bottles during production, recorded the measurements and plotted them as shown in Figure 13.14. This control chart reveals that only with difficulty could the process average be kept in control. Occasional operator interventions are required. Also the process range is moving towards (and once exceeding) the upper control limit. The process also seems to be becoming more variable. (After investigation, it was discovered that, because of faulty maintenance of the line, skin cream was occasionally contaminating the part of the line that fitted the cap, resulting in erratic tightening of the caps.)

Control charts for attributes

Attributes have only two states – 'right' or 'wrong', for example – so the statistic calculated is the proportion of wrongs (p) in a sample. (This statistic follows a binomial distribution.) Control charts using p are called 'p-charts'. When calculating control limits, the population mean ($\overline{p}$) (the actual, normal or expected proportion of 'defectives') may not be known. Who knows, for example, the actual number of city commuters who are dissatisfied with their journey time? In such cases, the population mean can be estimated from the average of the proportion of 'defectives' ($\overline{p}$), from m samples each of n items, where m should be at least 30 and n should be at least 100:

$$\overline{p} = \frac{p^1 + p^2 + p^3 \ldots p^n}{m}$$

One standard deviation can then be estimated from:

$$\sqrt{\frac{\bar{p}(1-\bar{p})}{n}}$$

The upper and lower control limits can then be set as:

$$\text{UCL} = \bar{p} + 3 \text{ standard deviations}$$

$$\text{LCL} = \bar{p} - 3 \text{ standard deviations}$$

Of course, the LCL cannot be negative, so when it is calculated to be so it should be rounded up to zero.

WORKED EXAMPLE

A credit card company deals with many hundreds of thousands of transactions every week. One of its measures of the quality of service it gives its customers is the dependability with which it mails customers' monthly accounts. The quality standard it sets itself is that accounts should be mailed within two days of the 'nominal post date', which is specified to the customer. Every week the company samples 1,000 customer accounts and records the percentage not mailed within the standard time. When the process is working normally, only 2 per cent of accounts are mailed outside the specified period, that is, 2 per cent are 'defective'.

Control limits for the process can be calculated as follows:

$$\text{Mean proportion defective, } \bar{p} = 0.02$$

$$\text{Sample size } n = 1{,}000$$

$$\text{Standard deviation } s = \sqrt{\frac{\bar{p}(1-\bar{p})}{n}}$$

$$= \sqrt{\frac{0.02(0.98)}{1{,}000}}$$

$$= 0.0044$$

With the control limits at $\bar{p} \pm 3s$:

$$\text{Upper control limit (UCL)} = 0.02 + 3(0.0044) = 0.0332$$

$$= 3.32\%$$

$$\text{and lower control limit (LCL)} = 0.02 - 3(0.0044) = 0.0068$$

$$= 0.68\%$$

Figure 13.14 shows the company's control chart for this measure of quality over the last few weeks, together with the calculated control limits. It also shows that the process is in control.

Sometimes it is more convenient to plot the actual number of defects (c) rather than the proportion (or percentage) of defectives, on what is known as a c-chart. This is very similar to the p-chart but the sample size must be constant and the process mean and control limits are calculated using the following formulae:

$$\text{Process mean } \bar{c} = \frac{c_1 + c_2 + c_3 \ldots c_m}{m}$$

$$\text{Control limits} = \bar{c} \pm 3\sqrt{c}$$

where c = number of defects
m = number of samples

Applying the principles

1. An animal park in Amsterdam has decided to sample 50 visitors each day (n) to see how many visitors are from overseas. The data in Table 13.7 below are for the last 7 days. If it decided to continue recording this data and plot it on a control chart for attributes, what should the upper and lower control limits be?
2. The manager of a sweet shop decides to sample batches of sweets to check that the weight is reasonably consistent. She takes 9 samples, each with 10 bags. The data in Table 13.8 below shows the average mean weight for each sample and the weight range. What control limits would a control chart for variables use?

Table 13.7 Number of visitors from overseas

1	7
2	8
3	12
4	5
5	5
6	4
7	8

Table 13.8 The average mean weight for each sample and the weight range

Sample	*Weight average in grams* $\bar{x}$	*Range* (R)
1	10	1.50
2	8	2.00
3	9	3.00
4	9	2.50
5	8	1.50
6	9	1.00
7	11	2.00
8	14	2.50
9	12	2.00

TAKING IT FURTHER

Dale, B. G. van der Wiele, T. and van Iwaarden, J. (2007) ***Managing Quality (5th edn), Wiley-Blackwell.*** *This is the fifth edition of a book that has long been one of the best respected texts in the area. A comprehensive and balanced guide to the area.*

Oakland, J.S. (2007) ***Statistical Process Control, Routledge.*** *Comprehensive, but not too daunting.*

Pande, P.S., Neuman, R.P. and Kavanagh, R.R. (2000) ***The Six Sigma Way, McGrawHill.*** *There are many books written by consultants for practicing managers on the now fashionable Six Sigma Approach. This one is readable and informative.*

Webber, L. and Wallace, M. (2007) ***Quality Control for Dummies, John Wiley & Sons.*** *Written for the practitioner, this is an easy guide to improving quality. It covers expert techniques for introducing quality methods.*

Yadav, N. (2017) ***Implementation of Statistical Process Control (SPC) for Manufacturing, LAP LAMBERT Academic Publishing.*** *An up to date treatment, but relatively technical and manufacturing focused.*

14 Risk and resilience

Introduction

All operations face risk and experience failure. Some risks of failure may emerge within the operation; through poor quality processes, capacity constraints, or short-term financial targets. Some risks emerge as a consequence of supply network design; for example, reliance on over-powerful or unreliable suppliers, supply complexity, lack of visibility, or uncertainty in demand. Other risks come from broader environmental (also called institutional) forces, such as economic downturns, political unrest and environmental disasters. In the face of such risks, a 'resilient' operation or process is one that can identify likely sources of risks, prevent failures occurring, minimise their effects and learn how to recover from them. In a world where failures hurt organisations more and more, managing risk is an increasingly vital task (see Figure 14.1).

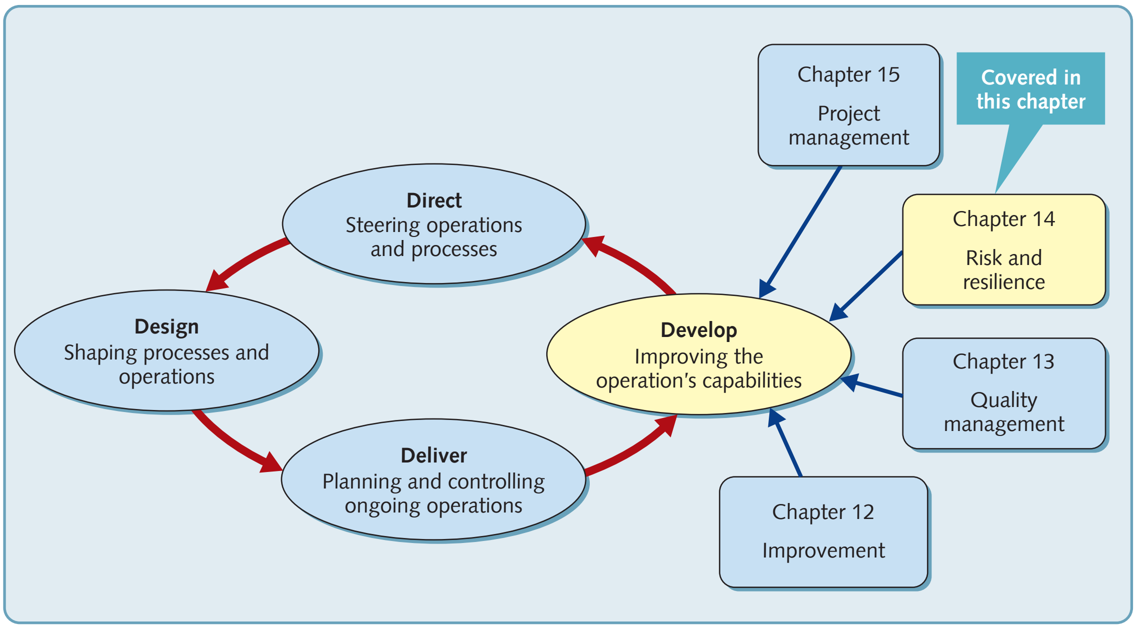

Figure 14.1 Risk is the negative consequences from events; resilience is preventing, mitigating and recovering from these events

EXECUTIVE SUMMARY

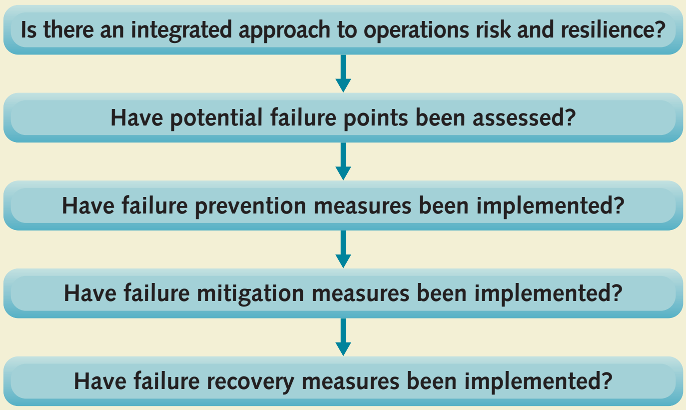

Is there an integrated approach to operations risk and resilience?

Risk is the potential for unwanted negative consequences from events. Resilience is the ability to prevent, withstand and recover from those events. Failure events can be categorised in terms of the seriousness of their impact and the likelihood of their occurrence. Relatively low-impact failures that happen relatively frequently are the province of quality management. Resilience involves attempting to reduce the combined effects of a failure occurring and the negative impact that it may have. Comprehensive risk and resilience management involves integrating four sets of activities, understanding and assessing the seriousness of the potential failures, preventing failures, minimising their negative consequences (called failure mitigation) and recovering from failure so as to reduce its impact.

Have potential failure points been assessed?

Managing risk begins by identifying potential sources of failure. These sources can be categorised into supply failures, failures that happen inside the operation (further categorised as human, organisational and technology failures), product/service design failures, customer failures and the failures caused by environmental (or institutional) forces, such as political upheaval, war, weather, crime, and so on. Understanding why such failures occur can be aided by post-failure analysis using accident investigation, traceability, complaint analysis, fault-tree analysis and other similar techniques. Judging the likelihood of failure may be relatively straightforward for some well-understood processes, but often has to be carried out on a subjective basis, which is rarely straightforward.

Have failure prevention measures been implemented?

Once the potential sources of failure have been identified, managers must examine ways to prevent or reduce the likelihood of failures. The main approaches to failure prevention involve designing out the possibility of failure at key points in a process; installing fail-safeing mechanisms that prevent the mistakes which can cause failure; maintaining processes so as to reduce the likelihood of failure; and providing extra but redundant resources that can provide back-up in the case of failure.

Have failure mitigation measures been implemented?

Even when a failure has occurred, its impact on the customer can, in many cases, be minimised through mitigation actions. Failure mitigation means isolating a failure from its negative consequences. There are various mitigation actions, including economic mitigation using insurance; risk sharing and hedging; spatial or temporal containment that prevents the failure spreading geographically or over time; loss reduction that removes whatever might be harmed by a failure; and substitution that involves providing substitute resources to work on a failure before it becomes serious.

Have failure recovery measures been implemented?

When failure occurs and cannot (or can only partially) be mitigated, operations managers must carefully consider alternative methods of recovery. Failure recovery is the set of actions that are taken to reduce the impact of failure once its negative effects have been experienced by the customer. Recovery needs to be planned and procedures put in place that can discover when failures have occurred, guide appropriate action to keep everyone informed, capture the lessons learnt from the failure and plan to absorb lessons into any future recovery.

DIAGNOSTIC QUESTION

Is there an integrated approach to operations risk and resilience?

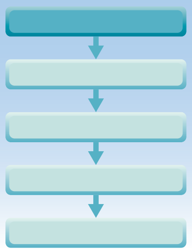

Managing operations and resilience is one of any practising operations manager's most important roles. Risk is the potential for unwanted negative consequences from some event. Resilience is the ability to prevent, withstand and recover from those events. Failures (things that have negative consequences) happen in all operations, but accepting that failures occur is not the same thing as tolerating or ignoring them. Operations do generally attempt to minimise both the likelihood of failure and the effect it will have, although the method of coping with failure will depend on how serious are its negative consequences, and how likely it is to occur. At the lower end of the scale, the whole area of quality management is concerned with identifying and reducing every small error in the creation and delivery of products and services. Other failures will have more impact on the operation, even if they do not occur very frequently. For example, a server failure can seriously affect service and therefore customers, which is why system reliability is such as important measure of performance for IT service providers. Finally, some failures, while much less likely, are so serious in terms of negative consequences that we class them as disasters. Examples may include major floods, earthquakes and hurricanes, wars and acts of terrorism, and financial collapses such as a stock market crash.

> **OPERATIONS PRINCIPLE**
> *Failure will always occur in operations, recognising this does not imply accepting or ignoring it.*

This chapter is concerned with all types of failure other than those with relatively minor consequences. This is illustrated in Figure 14.2. Some of these failures are irritating, but relatively unimportant, especially those close to the bottom left-hand corner of the matrix in Figure 14.2.

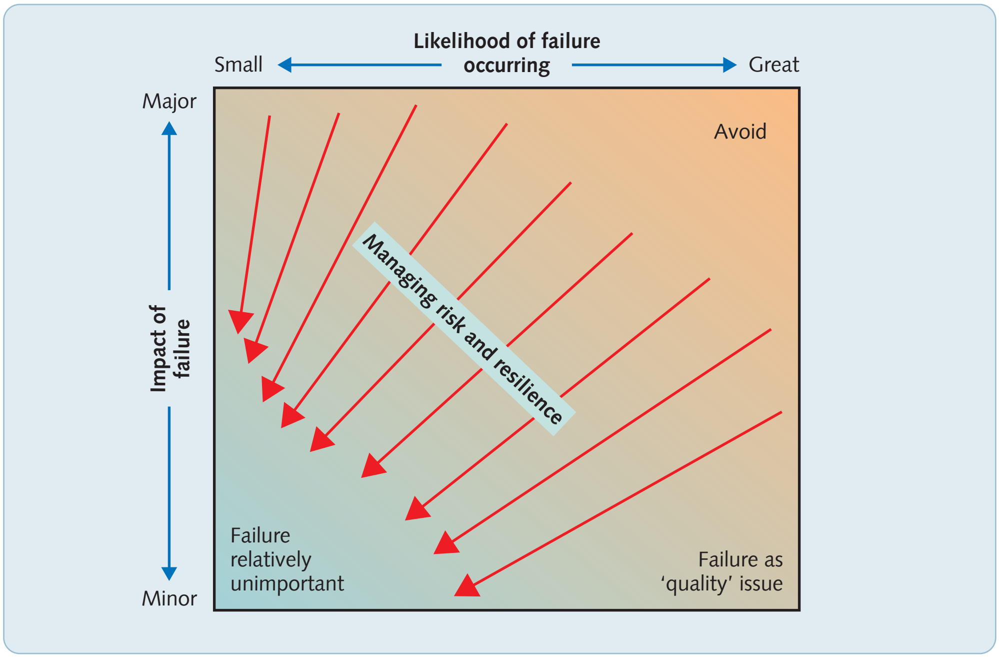

Figure 14.2 How failure is managed depends on its likelihood of occurrence and the negative consequence of failure

Other failures, especially those close to the top right-hand corner of the matrix, are normally avoided by all businesses because embracing such risks would be clearly foolish. In-between these two extremes is where most operations-related risks occur. In this chapter, we treat various aspects of these types of failure, and in particular how they can be moved in the direction of the arrows in Figure 14.2.

EXAMPLES

So many risks[1]

Close to the start of all the other chapters in this text, two examples are presented that illustrate some of the issues covered. This chapter is different. To illustrate the range of operations-related risks that businesses face, here are six examples.

The central bank of Bangladesh heist

In 2016, a cyber-attack was mounted on the central bank of Bangladesh. $850 million in transactions were prevented from being processed and the bank lost $81million. It appears that hackers had used stolen credentials to transfer nearly $1 billion from Bangladesh Bank's account at the Federal Reserve Bank of New York to bank accounts in the Philippines, Sri Lanka and other parts of Asia. A Bangladesh government-appointed panel investigating the cyber-heist found the hackers may have manipulated loopholes in the bank's online security when their technical staff connected their local money transfer system with the SWIFT international payments network. SWIFT is a consortium that operates a computer worldwide network for communication between banks. The hackers used sophisticated techniques and, according to investigators, displayed both technical skills and a deep knowledge of the operations infrastructure used by Bangladesh Bank to interface with SWIFT. Later, the panel found five officials at the bank guilty of negligence and carelessness, but as to the perpetrators themselves? Suggestions include cybercriminals from organised crime syndicates, insiders, ex-employees with a grudge and the North Korean government. But whoever carried out the attack, their task was made easier by a lack of adequate security.

Cadbury's salmonella outbreak

In June 2007, Cadbury, the confectionary company, was fined £1 million plus costs for breaching food safety laws when a national salmonella outbreak infected 42 people, including children aged under 10, who became ill with a rare strain of Salmonella Montevideo. '*I regard this as a serious case of negligence*', the judge said. One prominent lawyer announced that '. . . *Despite Cadbury's attempts to play down this significant fine, make no mistake it was intended to hurt and is one of the largest of its kind to date*.' The company apologised offering its '*sincere regrets*' to those affected, and pleaded guilty to nine food safety offences. But Cadbury admitted that it failed to notify the authorities of positive tests for salmonella as soon as they were known within the company. Cadbury, through its layers, said: '*Negligence we admit, but we certainly do not admit that this was done deliberately to save money and nor is there any evidence to support that conclusion*.' The judge said Cadbury had accepted that a new testing system, originally introduced to improve safety, was a '*distinct departure from previous practice'*, and was '*badly flawed and wrong*'. In a statement Cadbury said: '*Mistakenly, we did not believe that there was a threat to health and thus any requirement to report the incident to the authorities – we accept that this approach was incorrect. The processes that led to this failure ceased from June last year and will never be reinstated*.' The company was not only hit by the fine and court costs, it had to bear the costs of recalling one million bars that may have been contaminated from sale, and face private litigation claims brought by its consumers who were affected. Cadbury said it lost around £30 million because of the recall and subsequent safety modifications, not including any private litigation claims.

Tesco, Findus and the horsemeat scandal

In 2013, several news organisations reported that horse DNA had been discovered in frozen beef burgers that were widely available for sale in British and Irish supermarkets. Subsequently, a long list of countries also reported findings of adulterated beef following the initial discovery, including Sweden, Poland, Italy, Germany, Greece, the Netherlands and Austria. Horsemeat is not harmful to health (it is eaten in many countries), but consumers in the UK and Ireland do not generally view it as acceptable. Further investigations revealed high levels of horsemeat in 'beef' products, and in some cases pig DNA, a taboo food for people in the Muslim and Jewish communities. A lack of sourcing traceability in the supply chain was blamed for the problem. The source of the adulterated meat seems to have been French company Comigel, who said that it came from their supplier, Spanghero, which was investigated on charges of aggravated fraud, resulting in the arrest of eight managers. In the UK, the press focused on two companies: Tesco, the UK's largest supermarket chain, which had sold some of the offending products; and Findus, a frozen food brand that had produced 'beef' lasagne product found by the UK's Food Standards Agency to contain 100 per cent horsemeat. A subsequent investigation revealed that the reputation of both Tesco and Findus has suffered. Nearly £300 million was wiped off the value of Tesco shares. Yet there was a significant difference in how the two companies' responded to the crisis. Tesco's response was particularly successful. It faced the crisis head-on, willing to take ownership of the problem, proactive in its communications, fast in responding and with a direct and clear message. Findus, by contrast, tended to keep quiet as the scandal unfolded.

G4S and its Olympic adventure

On 21 May 2013 Nick Buckles, stepped down as chief executive of G4S, a leading global security company. Despite share growth and expansion of operations, for many, Mr Buckles' leadership will be remembered for the company's failure to fulfil its obligations to provide security to the 2012 London Olympic Games. A year earlier, G4S had been given the contract to provide training and management of the 10,000 security personnel that would be required for the games. By July 2012, media reports suggested that G4S were struggling to recruit and train staff. Worse still, a 'whistle-blower' claimed that security-vetting procedures were being overlooked and some employees were 'self-certifying'. Though initially denying these allegations, G4S were eventually forced to admit that it would not be able to deliver the number of security personnel that it had promised. (More embarrassingly, it later emerged that the problem had been flagged-up internally months earlier.) When the scale of the problem did become clear, the British Government announced that it would have to call upon an extra 3,500 troops, who would be used for security duties. It became clear that insufficient planning had taken place, nor was the implementation of the project tightly controlled. Later, G4S ruled out bidding for security contracts at the 2016 Olympics in Brazil because of 'the reputational damage that would be done to the company if the fiasco of the London contract was repeated.' By February 2013, G4S has announced that its losses would be even bigger than expected, standing at £70 million.

Volkswagen 'dieselgate' scandal

In 2015, the US Environmental Protection Agency (EPA) found that Volkswagen (VW) had been installing a piece of software into computers on its cars that falsified emissions data on its vehicles with diesel engines. The software (a so-called 'defeat device') could recognise when a car is being tested so that it could retune the engine's performance to limit nitrogen oxide emissions. After any test, when the car returned to normal road conditions, the level of such emissions increased sharply. The number of cars with the devices that enabled them to falsify their emission levels was estimated at 11 million cars worldwide. As the scale of the scandal became evident, VW's America boss Michael Horn said, '*We've totally screwed up*', and Martin Winterkorn, the group's then chief executive, said that VW had '*broken the trust of our customers and the public*' (he later resigned because of the scandal). Over a year after the news broke, the US Department of Justice announced that Volkswagen was to pay $4.3 billion under a plea deal with US authorities. This was in addition to a $15 billion civil settlement with car owners and environmental authorities in America, where VW had agreed to buy back some of the affected vehicles. However, appearing before a UK parliamentary committee in 2017, the head of Volkswagen UK, Paul Willis, said his company has '*not misled anybody*', nor does it owe any compensation to British drivers. '*We did not falsify information and we completely refute that we have misled anybody*', he said. '*You cannot compare the United States and Europe, because the regulations are very different. In the US, technical fixes do not get the cars back to compliance with the regulations, but in Europe they do, and they can be done in 20 or 30 minutes. There is no evidence that there has been any degradation in the resale value of these vehicles – so there is no loss. There is no legal basis for compensation*.'

The launch of Heathrow Terminal 5

Heathrow Terminal 5, which is used for British Airways (BA) flights, is now one of the best European airport terminals in operation, but its launch was a catastrophe. Anyone who was traveling through the terminal when it opened will never forget the chaos. It was, what BA's boss Willie Walsh called, with magnificent understatement, 'not the company's finest hour'. The mess made news around the world and was seen by many as one of the most public failures of basic operations management in aviation history. Two hundred flights were cancelled in its first three days, some customers were still without their luggage after three weeks, and it needed an extra 400 volunteer staff to wade through the backlog of late baggage. It was all caused by several interrelated problems. Press reports initially blamed glitches with the new baggage handling system. And while the baggage handling system did have problems, it was not the only, or possibly the main, problem. The system had worked until it became clogged with bags that were overwhelming BA's handlers loading them onto the aircraft. Partly this was because staff were not sufficiently familiar with the new operating processes, but handling staff had also suffered delays getting to their new (unfamiliar) work areas, negotiating (new) security checks and finding (again, new) car parking spaces. To make things worse, once staff were airside they had problems logging-in. The total cost of the crisis to BA was estimated to be £16 million. Shortly after the opening, BA announced that two of its most senior executives, its director of operations and its director of customer services would leave the company.

What do they have in common?

In terms of the causes of these operations-related 'failures', there seems not to be a lot that connects these examples. That is the point. Operations can fail for a whole variety of reasons. Yet there are some common issues that the examples demonstrate. The first is the damaging effects of major failures within operations and across supply networks. The impact of such failures can often be more significant than initially realised. Second, the consequences for an organisation's profits, reputation and share value can be very serious. Third, at least some aspects of the causes of the ensuing failure were, in fact, known within the companies but were largely ignored. The root causes of all the examples, with the possible exception of the central bank of Bangladesh heist, were internal to the organisations concerned (and one could argue that it was internal failure that allowed external criminals to hit the central bank of Bangladesh). Fourth, the examples highlight the role that media, government and other external stakeholders play during the failure and recovery process. As such, understanding how to manage and communicate with external stakeholders is critical in mitigating the negative consequences of failures. In the long-term, the reputational risks associated with operations-related failures can be even more significant than the immediate financial costs.

Assessment, prevention, mitigation and recovery

Obviously, some businesses function in a riskier environment than others. Those operations with a high likelihood of failure and/or serious consequences deriving from that failure will need to give it more attention. But operations and process resilience is relevant to all organisations. They all must give attention to the four sets of activities that will determine their resilience. The first is concerned with understanding what failures could potentially occur in the operation and assessing their seriousness. The second is to examine ways of preventing failures occurring. The third is to minimise the negative consequences of failure (called 'mitigation'). The final task is to devise plans and procedures that will help the operation to recover from failures when they do occur. The remainder of this chapter deals with these four tasks (see Figure 14.3).

OPERATIONS PRINCIPLE

Risk is governed by the effectiveness of failure assessment, failure prevention, mitigation and recovery.

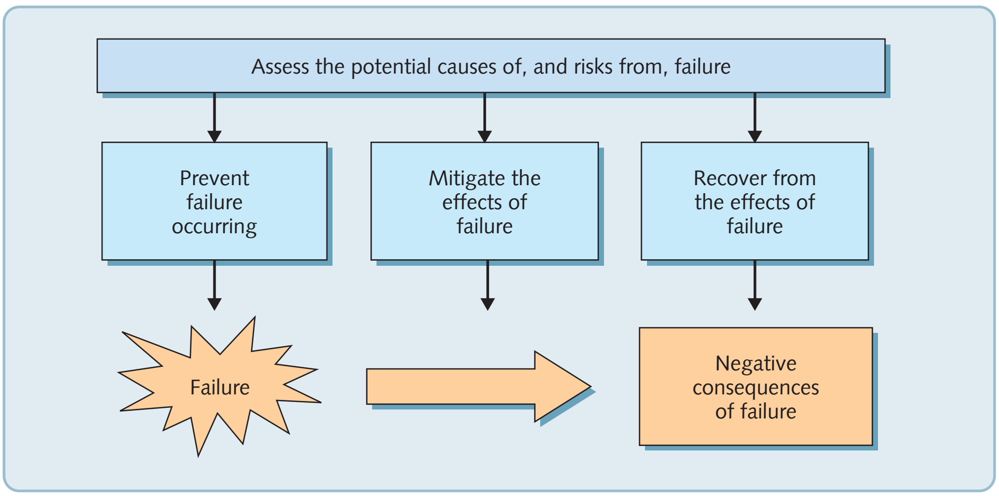

Figure 14.3 **Managing operations and process risk involves assessing the potential causes of, and risks from failure, failure prevention, mitigating the negative consequences of failure, and failure recovery**

DIAGNOSTIC QUESTION

Have potential failure points been assessed?

A prerequisite to achieving operations and process resilience is to understand where failure might occur and what the consequences of failure might be, by assessing causes of failure – often, it is the 'failure to understand failure' that leads to excessive disruption. Assessing potential failure points can be done in four ways. First, operations managers can identify different potential failure sources (supply, human, organisational, technology, product/service design, customer and environmental) to understand the likelihood of failure *before* events occur. Second, managers can use post-failure analysis techniques, such as accident investigations, failure traceability, complaint analysis and fault-tree analysis, to understand why failures occurred *after* the event. Third, they can combine the previous two approaches to calculate the likelihood of different failures. Finally, information on the consequences, probability and likelihood of detection can be used to assess the relative priorities of different threats, typically using failure mode and effect analysis (FMEA) (see later).

OPERATIONS PRINCIPLE

A 'failure to understand failure' is the root cause of a lack of resilience.

Identify the potential causes of failure

The causes of some failure are purely random, like lightning strikes, and are difficult, if not impossible, to predict. However, the vast majority of failures are not like this. They are caused by something that could have been avoided, which is why, as a minimum starting point a simple checklist of failure causes is useful. Figure 14.4 illustrates how this might be done. Here, failure sources are classified as: (1) failures of supply, internal failures such as those deriving from human organisational and technological sources; (2) failures deriving from the design of products and services; (3) failures deriving from customer failures; and (4) general environmental (or institutional) failures.

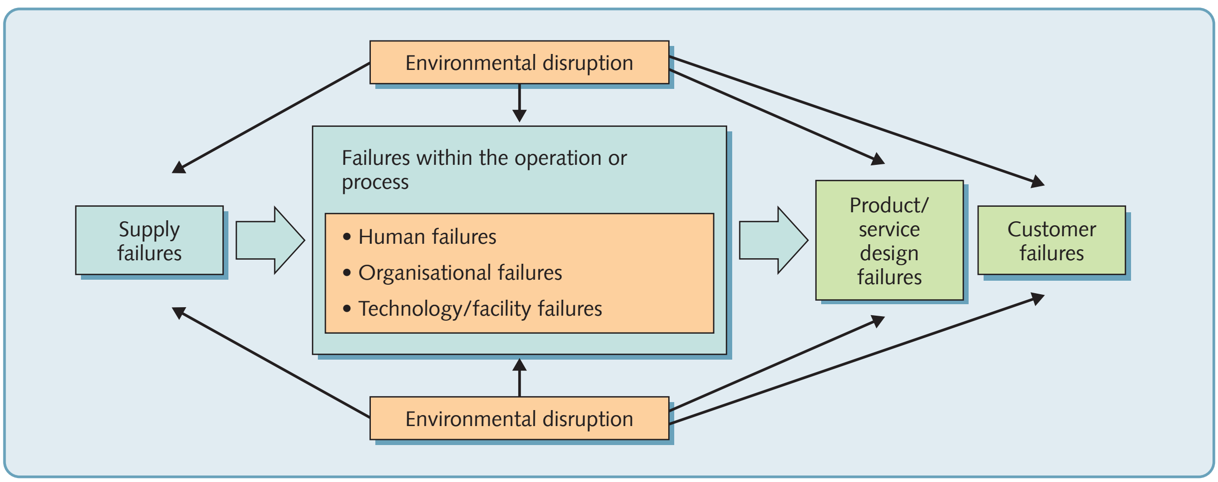

Figure 14.4 The sources of potential failure in operations

Supply failure

Supply failure means any failure in the timing or quality of products and services delivered into an operation; for example, suppliers delivering the wrong or faulty components, outsourced call centres suffering a telecoms failure, disruption to power supplies, and so on. The more an operation relies on suppliers for materials or services, the more it is at risk from failure caused by missing or sub-standard inputs. It is an important source of failure because of the increasing dependence on out-sourced activities in many industries, and the emphasis on keeping supply chains 'lean' in order to cut costs. Other factors have also increased exposure to supply failure in recent years. For example, the rise of global sourcing usually means that parts are shipped around the world on their journey through the supply chain. Microchips manufactured in Taiwan could be assembled to printed circuit boards in Shanghai that are then finally assembled into a computer in Ireland and delivered to the United States. At the same time, several industries are suffering increased volatility and uncertainty in demand. This may be as a result of market fragmentation, where small customer segments have to be accommodated. It may be as a result of faster changes in products, services and the supply base. But the result is to make demand forecasting extremely challenging. Perhaps most significantly there tends to be far less inventory in supply chains that could buffer interruptions to supply.

Human failures

There are two broad types of human failure. The first is where key personnel leave, become ill, die, or in some way cannot fulfil their role. The second is where people are actively doing their job but are making mistakes. Understanding risk in the first type of failure involves identifying the key people without whom operations would struggle to operate effectively. These are not always the most senior individuals, but rather those fulfilling crucial roles that require special skills or tacit knowledge. Human failure through 'mistakes' also comes in two types: errors and violations. 'Errors' are mistakes in judgement: with hindsight, a person should have done something different. For example, if the manager of a sports stadium fails to anticipate dangerous crowding during a championship event, this is an error of judgement. 'Violations' are acts that are clearly contrary to defined operating procedure. For example, if a maintenance engineer fails to clean a filter in the prescribed manner, it is eventually likely to cause failure. Catastrophic failures are often caused by a combination of errors and violations. For instance, one kind of accident, where an aircraft appears to be under control and yet still flies into the ground, is very rare (once in two million flights).[2] For this type of failure to occur, first, the pilot has to be flying at the wrong altitude (error). Second, the co-pilot would have to fail to cross-check the altitude (violation). Third, air traffic controllers would have to miss the fact that the plane was at the wrong altitude (error). Finally, the pilot would have to ignore the ground proximity warning alarm in the aircraft, which can be prone to giving false alarms (violation).

Organisational failure

Organisational failure is usually taken to mean systemic failures of procedures and processes and failures that derive from a business's organisational structure and culture. This is a huge potential source of failure and includes almost all operations and process management. In particular, failure in the design of processes (such as bottlenecks causing system overloading) and failures in the resourcing of processes (such as insufficient capacity being provided at peak times) need to be investigated. But there are also many other procedures and processes within an organisation that can make failure more likely. For example, remuneration policy may motivate staff to work in a way that, although increasing the financial performance of the organisation, also increases its susceptibility to failure. Examples of this can range from sales people being so incentivised that they make promises to customers that cannot be fulfilled, through to investment bankers being more concerned with profit (up-side risk) than the risks of financial overexposure (down-side risk). This type of risk can derive from an organisational culture that minimises consideration of risk, or it may come from a lack of clarity in reporting relationships.

Technology / facilities failures

By 'technology and facilities' we mean all the IT systems, machines, equipment and buildings of an operation. All are liable to failure, or breakdown. The failure may be only partial, for example a machine that has an intermittent fault. Alternatively, it can be what we normally regard as a 'breakdown' – a total and sudden cessation of operation. Either way, its effects could bring a large part of the operation to a halt. For example, a computer failure in a supermarket chain could paralyse several large stores until it is fixed. Such IT failures have occurred for several leading banks over recent years. The root cause of these failures can often be traced back to inadequate maintenance of exceptionally complex information systems, an issue we shall turn to shortly in this chapter. In other cases, cyber-attacks on corporate information systems create a new, or at least increased, threat to operations.

Product / service design failures

Products and services often look great in design, but only when they have to cope with real demand do inadequacies become evident. A recent example of this is Heathrow's Terminal 5, which reported that it had not been able to replace a single light bulb in five years since its opening. The reason? When designing this iconic structure, no one thought to examine how its 120,000 light fittings would be maintained when they reached the end of their life. After several failed attempts at finding a solution, Heathrow Airport Holdings announced they would be hiring the services of a specialist high-level rope-work company to carry out the work – they declined to state what the cost of such skilled operators would be! Of course, during the design process, potential risk of failure should have been identified and 'designed out'. But one only has to look at the number of 'product recalls' or service failures to understand that design failures are far from uncommon. Sometimes this is the result of a trade-off between fast time-to-market performance and the risk of the product or service failing in operation. And, while no reputable business would deliberately market flawed products or services, equally most businesses cannot delay a product or service launch indefinitely to eliminate every single risk of failure.

Customer failures

Not all failures are (directly) caused by the operation or its suppliers. Customers may 'fail' in that they misuse products and services. For example, an IT system might have been well designed, yet the user could treat it in a way that causes it to fail. Customers are not 'always right'; they can be inattentive and incompetent. However, merely complaining about customers is unlikely to reduce the chances of this type of failure occurring. Most organisations will accept that they have a responsibility to educate and train customers, and to design their products and services so as to minimise the chances of failure.

Environmental disruption

Environmental disruption includes all the causes of failure that lie outside an operation's direct influence. Typically, such disasters include major political upheaval, hurricanes, floods, earthquakes, temperature extremes, fire, corporate crime, theft, fraud, sabotage, terrorism, other security attacks, and the contamination of products or processes. This source of potential failure has risen to near the top of many firms' agenda due to a series of major events over recent years. As operations become increasingly integrated (and increasingly dependent on integrated technologies such as information technologies), businesses are more aware of the critical events and malfunctions that have the potential to interrupt normal business activity and even stop the entire organisation.

EXAMPLE

Reputational risk from the supply chain[3]

The risks that are present within an operation can be difficult to manage – those in the operation's wider supply network can be even worse. This is especially true when the risk is to a firm's reputation for ethical behaviour. So when Human Rights Watch (a non-profit, nongovernmental human rights organisation) issued a report that complained about forced overtime and underage labour at Cambodian factories that supply high street stores Marks and Spencer, Gap, H&M and Adidas, it caused some consternation. The report claimed that Cambodian factories supplying the stores were mistreating their workers, forcing them to work overtime, firing them when pregnant and using underage labour. Many of the 73 factories that were investigated for the repot were subcontractors to primary suppliers, making it much harder, said the report, to hold the stores selling final products to account. However, some commented on the lack of specific identifiable cases in the report. '*Because of concerns for the job security of the workers involved, we do not provide the names of the factories implicated*', the report said. '*We recognise that this makes it more difficult for the brands to respond to specific cases, but hope that this will encourage them to address the broader issues*.' The report also described how the various high street brands had responded to the criticisms, choosing Marks and Spencer for particular criticism, mainly because of its secrecy about suppliers and because they '*did not respond to the substantive concerns Human Rights Watch raised with them around subcontracting*'. However, Marks and Spencer said that they regularly audited suppliers and did not allow unauthorised subcontracting. Nor had they been presented with any evidence to support the claims. '*All our suppliers must adhere to our strict ethical standards as a condition of working with us. All our supplier factories are audited regularly by third party, independent auditors and are visited by compliance managers*', said a spokesperson.

Post-failure analysis

While sources of failure can often be identified in advance of their occurrence, it is also valuable to use previous failures to learn about sources of potential risk. This activity is called post-failure analysis. Some techniques for this were described as 'improvement techniques' in Chapter 12. Others include the following.

Accident investigation

Large-scale national disasters like oil tanker spillages and aeroplane accidents are usually investigated using accident investigation, where specifically trained staff analyse the causes of the accident. In many senses, the reason so much attention goes into examining these kinds of failures after the event is not only because of the damaging consequences of failure, but also because their infrequency makes it relatively hard to identify new sources of risks in advance of an event.

Failure traceability

Some businesses (either by choice or because of a legal requirement) adopt traceability procedures to ensure that all their failures (such as contaminated food products) are traceable. Any failures can be traced back to the process that produced them, the components from which they were produced, or the suppliers who provided them. Bio-tagging in pharmaceuticals is one example of this approach. If a medical product is found, or suspected, to have a problem, all other products in the batch can be recalled.

Complaint analysis

Complaints (and compliments) are a potentially valuable source for detecting the sources of failures. The prime function of complaint analysis involves analysing the number and 'content' of complaints over time to understand better the nature of the failure, as the customer perceives it. Two key advantages of complaints are that they come unsolicited and also they are often very timely pieces of information that can pinpoint problems quickly. However, managers should be aware that for every customer who does complain, there may be many who do not.

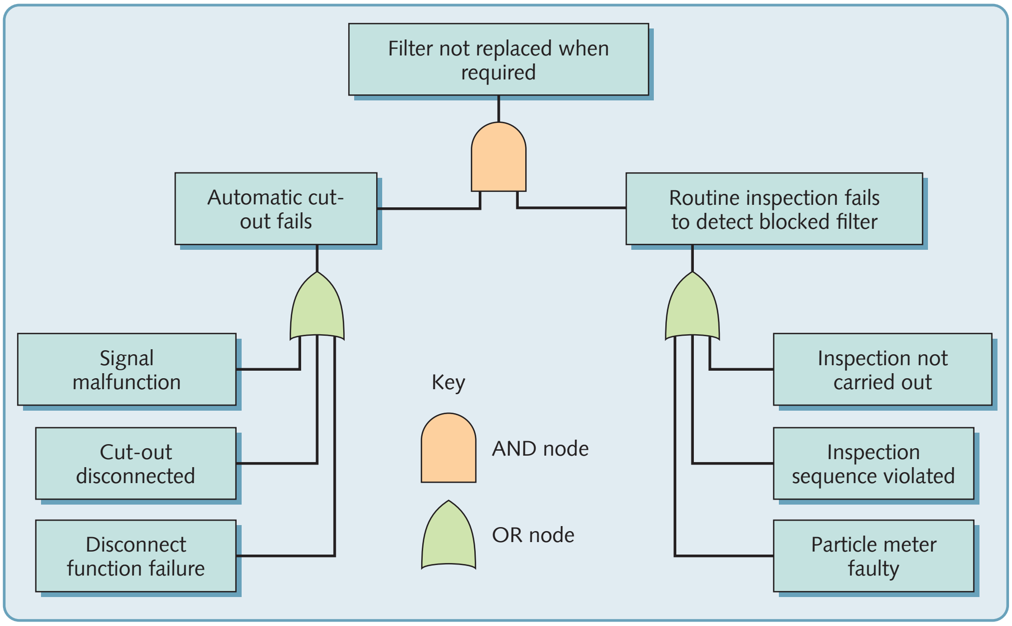

Figure 14.5 **Fault-tree analysis for failure to replace filter when required**

Fault-tree analysis

This is a logical procedure that starts with a failure or a potential failure and works backwards to identify all the possible causes and therefore the origins of that failure. Fault-tree analysis is made up of branches connected by two types of nodes: AND nodes and OR nodes. The branches below an AND node all need to occur for the event above the node to occur. Only one of the branches below an OR node needs to occur for the event above the node to occur. Figure 14.5 shows a simple tree identifying the possible reasons for a filter in a heating system not being replaced when it should have been.

Likelihood of failure

The difficulty of estimating the chance of a failure occurring varies greatly. Some failures are the result of well-understood phenomena. A combination of rational causal analysis and historical performance data can lead to a relatively accurate estimate of failure occurring. For example, a mechanical component in an air conditioning unit installed within an airport departure lounge may fail between 18 and 24 months of its installation in 95 per cent of cases. Other types of failure are far more difficult to predict. The chances of a fire in a bank's outsourced call centre are (hopefully) low, but how low? There will be some data concerning fire hazards in this type of operation, and one may insist on regular hazard inspection reports from the supplier's insurance providers, but the estimated probability of failure will be both low and subjective.

Assessing operational risks is not always easy. Even businesses that assess risks as part of their service can get it wrong. In what was one of the more ironic incidents, a worker at Aviva, the insurance company, reached an out-of-court compensation settlement after she fell over a pile of accident claim forms.

'Objective' estimates

Estimates of failure based on historical performance can be measured in several ways, including:

- *failure rates* – how often a failure occurs
- *reliability* – the chances of a failure occurring

- *availability* – the amount of available useful operating time.

'Failure rate' and 'reliability' are different ways of measuring the same thing: the propensity of an operation, or part of an operation, to fail. Availability is one measure of the consequences of failure in the operation. Sometimes failure is a function of time. For example, the probability of an electric lamp failing is relatively high when it is first used, but if it survives this initial stage, it could still fail at any point and the longer it survives, the more likely its failure becomes. Most physical parts of an operation behave in a similar manner. The curve which describes failure probability of this type is called the bath-tub curve. It comprises three distinct stages:

- the 'infant-mortality' or 'early-life' stage where early failures occur, caused by defective parts or improper use
- the 'normal-life' stage when the failure rate is usually low and reasonably constant, and caused by normal random factors
- the 'wear-out' stage when the failure rate increases as the part approaches the end of its working life, and failure is caused by the ageing and deterioration of parts.

Figure 14.6 illustrates three bath-tub curves with slightly different characteristics. Curve A shows a part of the operation which has a high initial infant-mortality failure but then a long, low-failure, normal-life followed by the gradually increasing likelihood of failure as it approaches wear-out. Curve B, while having the same stages, is far less predictable. The distinction between the three stages is less clear, with infant-mortality failure subsiding only slowly and a gradually increasing chance of wear-out failure. Failure of the type shown in curve B is far more difficult to manage in a planned manner. The failure of operations which rely more on human resources than on technology, such as some services, can be closer to curve C of Figure 14.6. They may be less susceptible to component wear-out but more so to staff complacency. Without review and regeneration, the service may become tedious and repetitive, and after an initial stage of failure reduction, as problems in the service are ironed out, there can be a long period of increasing failure.

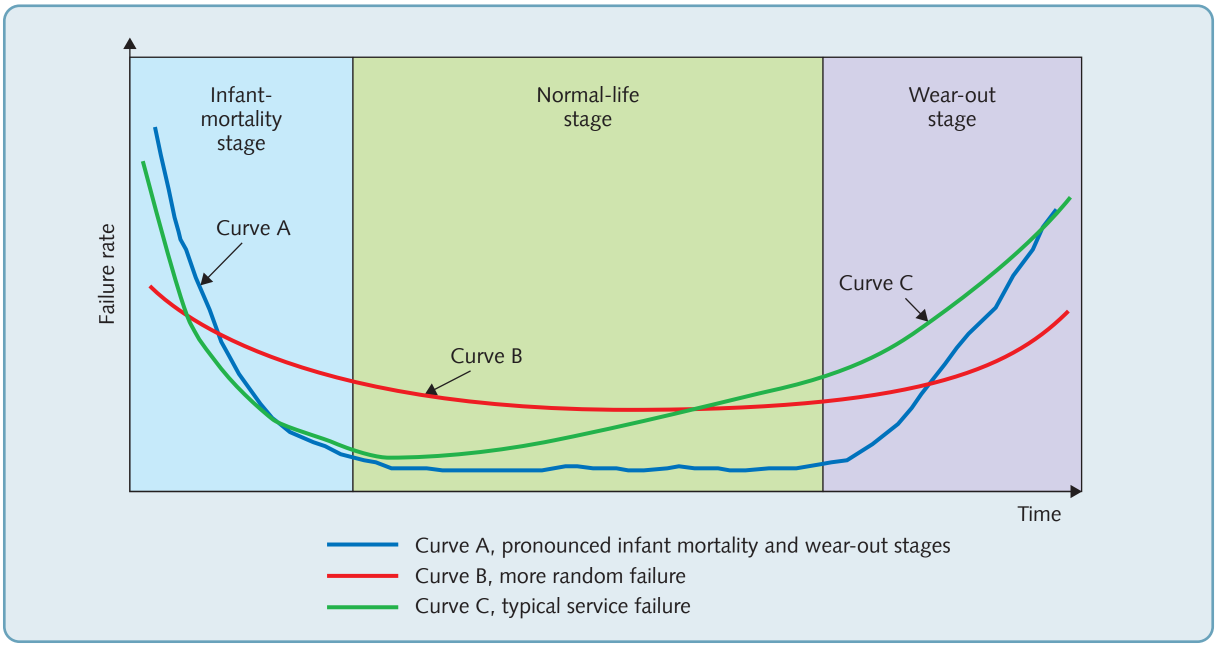

Figure 14.6 Bath-tub curves for three types of process

'Subjective' estimates

Failure assessment, even for subjective risks, is increasingly a formal exercise that is carried out using standard frameworks, often prompted by health and safety concerns, environmental regulations, and so on. These frameworks are similar to the formal quality inspection methods associated with quality standards like ISO 9000 that often implicitly assume unbiased objectivity. However, individual attitudes to risk are complex and subject to a wide variety of influences. In fact, many studies have demonstrated that people are generally very poor at making risk-related judgements. Consider the success of state and national lotteries. The chances of winning, in nearly every case, are extraordinarily low, and the costs of playing sufficiently significant to make the financial value of the investment entirely negative. If a player has to drive their car in order to purchase a ticket, they may be more likely to be killed or seriously injured en route than they are to win the top prize. But, although people do not always make rational decisions concerning the chances of failure, this does not mean abandoning the attempt. However, it does mean that one must understand the limits to overly rational approaches to failure estimation; for example, how people tend to pay too much attention to dramatic low-probability events and overlook routine events.

OPERATIONS PRINCIPLE

Subjective estimates of failure probability are better than no estimates at all.

EXAMPLE

The rise of the MicroMort[4]

Two authorities on risk, Michael Blastand and David Spiegelhalter, try to untangle the nature of risk and probability through the concept of the 'MicroMort'. The MicroMort is defined as a one-in-a-million chance of death and its use allows for some interesting comparisons to be made. For example, the chance of dying on a return motorbike trip from Edinburgh to London in the UK is around 120 MicroMorts (i.e. a 120-in-a-million chance). This is comparable to the risk faced by mothers giving birth in Britain or a soldier spending 2.5 days during the most dangerous period of the recent Afghanistan conflict. The concept also allows us to examine how risks have changed over time. For example, aircrews involved in bombing raids during the Second World War were 'exposed' to 25,000 MicroMorts – a depressing 2.5 per cent chance of death – per mission. Soldiers in Afghanistan in the recent conflict faced [just] 47 Micromorts – or 0.0047 per cent chance of death – per day on the ground. What is perhaps most interesting about the work is that it explores just how irrational humans are when attempting to calculate risk. Typically high-impact / low-probability risk is over-emphasised, while low-impact / high-probability risk is under emphasised. For organisations, appreciating the innate irrationality and cognition biases of most of its managers and employees can be an important step towards developing better assessments of various potential sources of failure.

Non-evident failure estimation

Not all failures are immediately evident. Small failures may accumulate for a while before they become evident, making objective and subjective estimation challenging. For example, purchasing managers who encounter difficulties in using an e-procurement system may simply circumvent the system and the failure points may only become evident when levels of non-compliance reach sufficient levels for senior management to notice. Likewise, within an automated production line, debris can accumulate. These may not cause an immediate failure, but could eventually lead to sudden and dramatic failure. Even when such failures are detected, they may not always receive the appropriate attention, either because there are inadequate failure identification systems, or

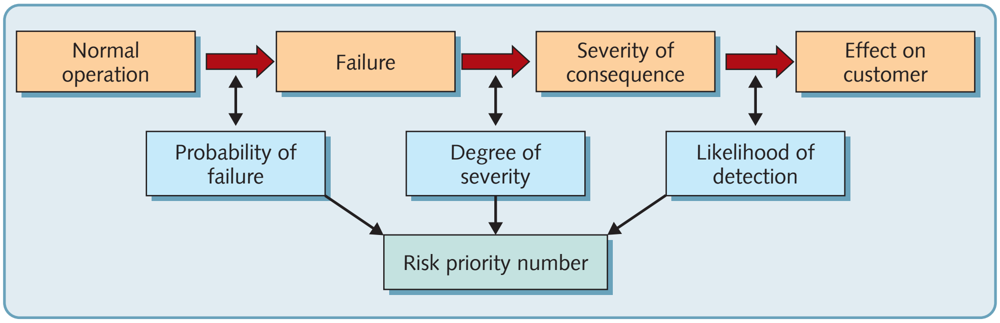

Figure 14.7 **Procedure for failure mode and effect analysis (FMEA)**

because there is insufficient managerial support or interest in making improvements. The mechanisms available to seek out failures in a proactive way include machine diagnostic checks, in-process checks, point-of-departure and phone interviews, and customer focus groups.

Failure mode and effect analysis

Having identified potential sources of failure (either in advance of an event of through post-failure analysis) and having then examined the likelihood of these failures occurring through some combination of objective and subjective analysis, managers can move to assigning relative priorities to risk. The most well-known approach for doing this is failure mode and effect analysis (FMEA). Its objective is to identify the factors that are critical to various types of failure as a means of identifying failures before they happen. It does this by providing a 'checklist' procedure built around three key questions for each possible cause of failure:

1. What is the likelihood that failure will occur?
2. What would the consequence of the failure be?
3. How likely is such a failure to be detected before it affects the customer?

Based on a quantitative evaluation of these three questions, a risk priority number (RPN) is calculated for each potential cause of failure. Corrective actions, aimed at preventing failure, are then applied to those causes whose RPN indicates that they warrant priority, see Figure 14.7.

DIAGNOSTIC QUESTION

Have failure prevention measures been implemented?

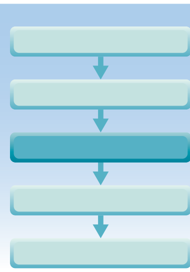

It has been said many times, 'prevention is better than cure'. This is why failure prevention is such an important part of operations and process resilience. There are a number of approaches to this, including designing out failure points, fail-safeing, maintenance and deploying redundant resources.

Designing out fail points

Process mapping, described in Chapter 6, can be used to 'engineer out' the potential fail points in various operations. For example, Figure 14.8 shows a process map for an automobile repair process. The stages in the process that are particularly prone to failure and the stages which are critical

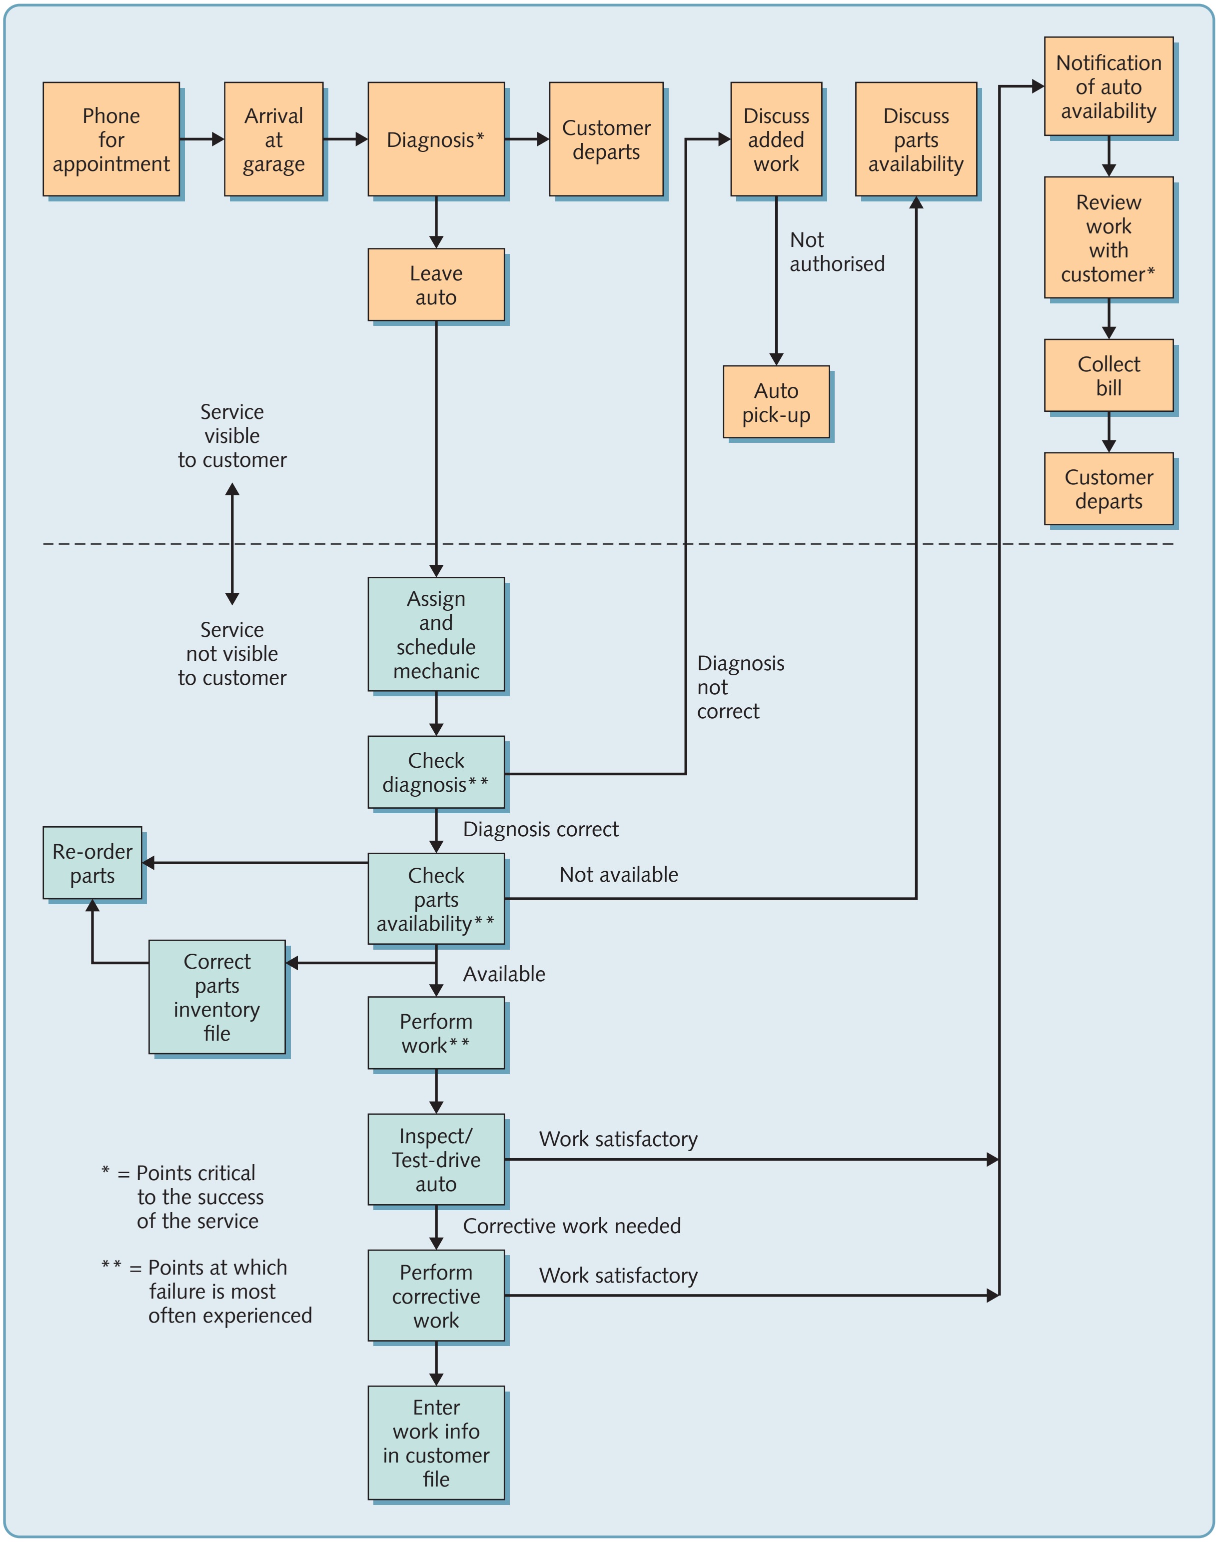

Figure 14.8 A process map for the auto repair process

to the success of the service have been marked. This will have been done by the staff of this operation metaphorically 'walking themselves through' the process and discussing each stage in turn.

Fail-safeing

The concept of fail-safeing has emerged since the introduction of Japanese methods of operations improvement. Called poka-yoke in Japan (from yokeru (to prevent) and poka (inadvertent errors)), the idea is based on the principle that human mistakes are to some extent inevitable. What is important is to prevent them becoming defects. Poka-yokes are simple (preferably inexpensive) devices or systems which are incorporated into a process to prevent inadvertent operator mistakes resulting in a defect. Typical poka-yokes include:

- trays used in hospitals with indentations shaped to each item needed for a surgical procedure – any item not back in place at the end of the procedure might have been left in the patient
- checklists which have to be filled in, either in preparation for, or on completion of, an activity, such as a maintenance checklist for a plane during turn-around
- gauges placed on machines through which a part has to pass in order to be loaded onto, or taken off, the machine – an incorrect size or orientation stops the process
- the locks on aircraft lavatory doors, which must be turned to switch the light on
- beepers on ATMs to ensure that customers remove their cards, or in cars to remind drivers to take their keys with them
- limit switches on machines that the machine to operate only if the part is positioned correctly
- height bars on amusement rides to ensure that customers do not exceed size limitations.

OPERATIONS PRINCIPLE

Simple methods of fail-safeing can often be the most cost effective.

Maintenance

While managers can try to design out failures and use fail-safe (poka-yoke) mechanisms to further reduce the likelihood of failures, operations and processes also need maintaining. Maintenance is the term typically used to cover the way operations and processes try to avoid failure by taking care of their physical facilities. It is particularly important when physical facilities play a central role in the operation, such as power stations, airlines and petrochemical refineries. There are a number of approaches to maintenance, including preventive maintenance, condition-based maintenance and total productive maintenance.

Preventive maintenance (PM)

This attempts to eliminate or reduce the chances of failure by regularly servicing facilities. For example, the engines of passenger aircraft are checked, cleaned and calibrated according to a regular schedule after a set number of flying hours.

Condition-based maintenance (CBM)

This attempts to perform maintenance only when the facilities require it. For example, continuous process equipment, such as that used in coating photographic paper, is run for long periods in order to achieve the high utilisation necessary for cost-effective production. Here CBM might involve continuously monitoring the vibrations, or some other characteristic of the line and then using the data to decide when to replace the bearings.

Total productive maintenance (TPM)

This combines aspects of PM and CBM approaches, but has a very strong focus on allowing people to take more responsibility for maintenance tasks. The approach looks to create a clear plan for all maintenance activities, including the level of preventive maintenance which is required for each piece of equipment, the standards for CBM and the respective responsibilities of operating staff and maintenance staff. For this to work, all staff must be trained in relevant maintenance skills and have all the skills to carry out their roles. Finally, TPM looks to use information from maintenance activities to steadily design out maintenance by considering what maintenance activities are a consequence of poor design, manufacturing, or installation.

How much maintenance?

Most operations plan their maintenance to include a level of regular preventive maintenance that gives a reasonably low but finite chance of breakdown. Usually, the more frequent the preventive maintenance episodes, the less are the chances of a breakdown. Infrequent preventive maintenance will cost little to provide but will result in a high likelihood (and therefore cost) of breakdown. Conversely, very frequent preventive maintenance will be expensive to provide but will reduce the cost of having to provide breakdown maintenance as in Figure 14.9(a). The total cost of maintenance appears to minimise at an 'optimum' level of preventive maintenance. However, this may not reflect reality. The cost of providing preventive maintenance in Figure 14.9(a) assumes that it is carried out by a separate set of people (skilled maintenance staff) whose time is scheduled and accounted for separately from the 'operators' of the facilities. In many operations, however, at least some preventive maintenance can be performed by the operators themselves (which reduces the cost of providing it) and at times which are convenient for the operation (which minimises the disruption to the operation). Furthermore, the cost of breakdowns could also be higher than is indicated in Figure 14.9(a) because unplanned downtime can take away stability from the operation, preventing it being able to improve itself. Put these two ideas together and the minimising total curve and maintenance cost curve look more like Figure 14.9(b). The emphasis is shifted towards using more preventive maintenance than is generally thought appropriate.

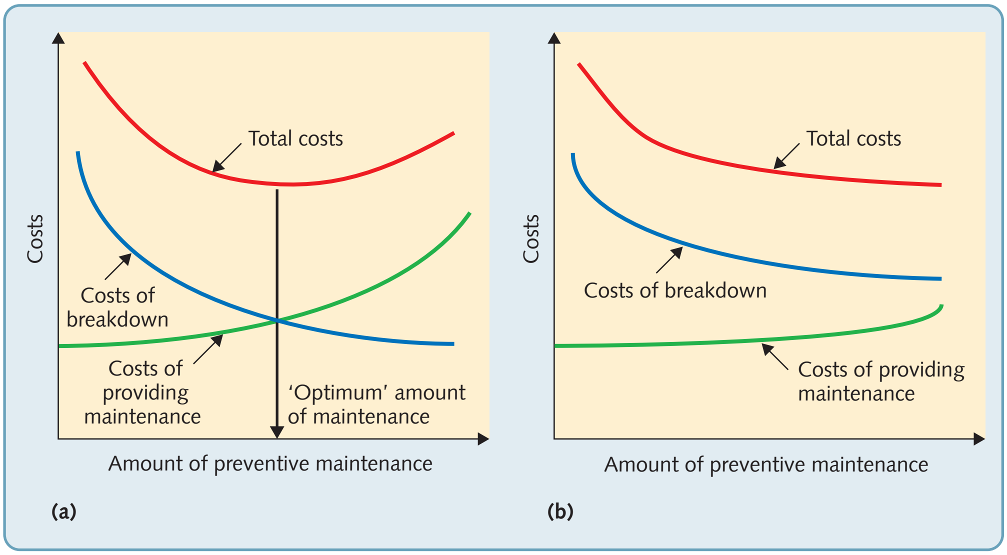

Figure 14.9 **Two views of maintenance costs (a) One model of the costs associated with preventive maintenance shows an optimum level of maintenance effort (b) If routine preventive maintenance tasks are carried out by operators and if the real cost of breakdowns is considered, the 'optimum' level of preventive maintenance shifts towards higher levels**

Redundancy

Building-in redundancy to an operation means having back-up processes or resources in case of failure. It can be an expensive solution to reduce the likelihood of the failure and is generally used when the breakdown could have a critical impact. Redundancy means doubling or even tripling some of the elements in a process so that these 'redundant' elements can come into action when a process fails. Nuclear power stations, hospitals and other public buildings have auxiliary or back-up electricity generators ready to operate in case the main electricity supply should fail. Some organisations also have 'back-up' staff held in reserve in case someone does not turn up for work, or is held up on one job and is unable to move on to the next. Spacecraft have several back-up computers on board that will not only monitor the main computer, but also act as a back-up in case of failure. Human bodies contain two of some organs – kidneys and eyes, for example – both of which are used in 'normal operation' but the body can cope with a failure in one of them. One response to the threat of large failures, such as terrorist activity, has been a rise in the number of companies (known as 'business continuity' providers) offering 'replacement office' operations, fully equipped with normal internet and telephone communications links, and often with access to a company's current management information. Should a customer's main operation be affected by a disaster, business can continue in the replacement facility within days or even hours.

The effect of redundancy can be calculated by the sum of the reliability of the original process component and the likelihood that the back-up component will both be needed and be working:

$$R_{a+b} = R_a + (R_b \times P_{(failure)})$$

where

R_{a+b} = reliability of component a with its back-up component b
R_a = reliability of a alone
R_b = reliability of back-up component b
$P_{(failure)}$ = the probability that component a will fail and therefore component b will be needed

So, for example, an e-auction service provider has two servers, one of which will come online only if the first server fails. If each server has a reliability of 0.9, the two working together (each with reliability = 0.9) will have a reliability of:

$$0.9 + [0.9 \times (1 - 0.9)] = 0.99$$

DIAGNOSTIC QUESTION

Have failure mitigation measures been implemented?

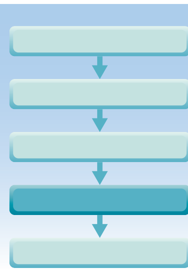

Failure mitigation means isolating a failure from its negative consequences. It is an admission that not all failures can be avoided. However, in some areas of operations, management relying on mitigation, rather than prevention, is unfashionable. For example, 'inspection' practices in quality management were based on the assumption that failures were inevitable and needed to be detected before they could cause harm. Modern total quality management places much more emphasis on prevention. Yet, in operations and process resilience, mitigation can be vital when used in conjunction with prevention in reducing overall risk.

Failure mitigation as a decision sequence

This whole topic involves managing under conditions of uncertainty. There may be uncertainty as to whether a failure has actually taken place at all. There almost certainly will be uncertainty about which courses of action will provide effective mitigation. There may even be uncertainty as to whether, what seems to have worked as a mitigation action, really has dealt with the problem. One way of thinking about mitigation is as a series of decisions under conditions of uncertainty. Doing so enables the use of formal decision analysis techniques such as decision trees, for example, which are illustrated in Figure 14.10. Here, an anomaly of some kind, which may or may not indicate that a failure has occurred, is detected. The first decision is whether to act to try and mitigate the supposed failure or, alternatively, wait until more information can be obtained. Even if mitigation is tried, it may or may not contain the failure. If not, then further action will be needed, which may or may not contain the failure, and so on. If more information is obtained prior to enacting mitigation, then the failure may or may not be confirmed. If mitigation is then tried, it may or may not work, and so on. Although the details of the specific mitigation actions will depend on circumstances, what is important in practical terms is that for all significant failures some kind of decision rules and mitigation planning has been established.

Failure mitigation actions

The nature of the action taken to mitigate failure will obviously depend on the nature of the failure. In most industries technical experts have established a classification of failure mitigation actions that are appropriate for the types of risk likely to be suffered. So, for example, in agriculture, government agencies and industry bodies have published mitigation strategies for such 'failures' as the outbreak of crop disease, contagious animal infections, and so on. Likewise, governments have different contingency plans in place to deal with the spread of major health risk such as the H1N1 and H7Np influenza viruses and an Ebola outbreak in West Africa. Such documents will outline the various mitigation actions that can be taken under different circumstances and detail exactly who are responsible for each action. Although these classifications tend to be industry specific, the following generic categorisation gives a flavour of the types of mitigation actions that may be generally applicable.

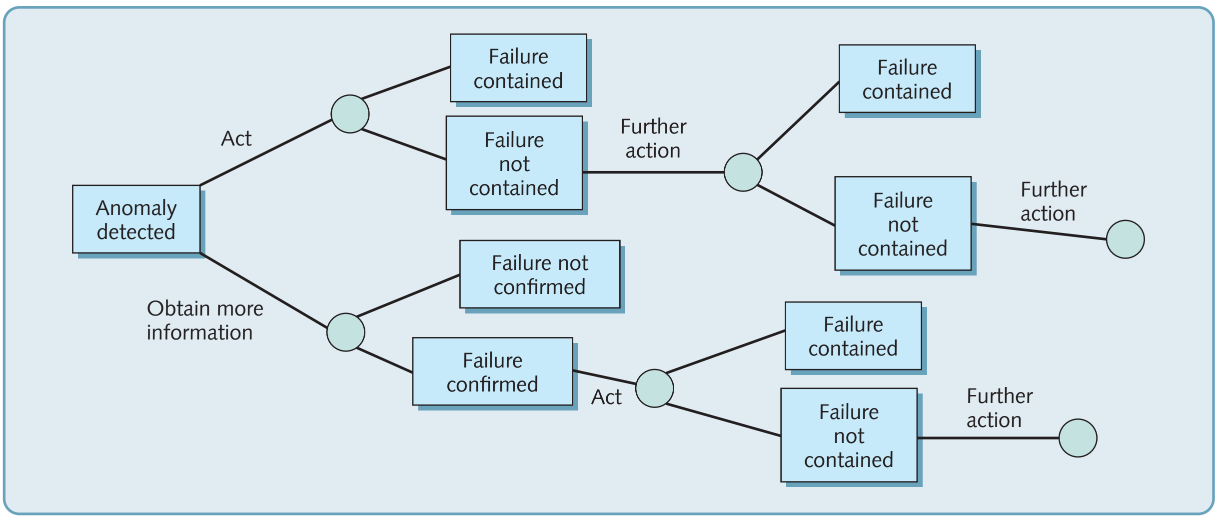

Figure 14.10 **A decision tree for mitigation when failure is not immediately obvious**

- **Mitigation planning** This is the activity of ensuring that all possible failure circumstances have been identified and the appropriate mitigation actions identified. It is the overarching activity that encompasses all subsequent mitigation actions, and may be described in the form of a decision tree or guide rules. Almost certainly there will be some form of escalation that will guide the extra mitigation effort, should early actions not prove successful. It is worth noting that mitigation planning, as well an overarching action, also provides mitigation action in its own right. For example, if mitigation planning has identified appropriate training, job design, emergency procedures, and so on, then the financial liability of a business for any losses should a failure occur will be reduced. Certainly businesses that have not planned adequately for failures will be more liable in law for any subsequent losses.
- **Economic mitigation** This includes actions such as insurance against losses from failure, spreading the financial consequences of failure and 'hedging' against failure. Insurance is the best known of these actions and is widely adopted, although ensuring appropriate insurance and effective claims management is a specialised skill in itself. Spreading the financial consequences of failure could involve, for example, spreading the equity holding in supply companies to reduce the financial consequences of such companies failing. Hedging involves creating a portfolio of ventures whose outcomes happen to be correlated so as to reduce total variability.
- **Containment (spatial)** This means stopping the failure physically spreading to affect other parts of an internal or external supply network. Preventing contaminated food from spreading through the supply chain, for example, will depend on real-time information systems that provide traceability data.
- **Containment (temporal)** This means containing the spread of a failure over time. It particularly applies when information about a failure or potential failure needs to be transmitted without undue delay. For example, systems that give advanced warning of hazardous weather such as snow storms must transmit the information to local agencies such as the police and road-clearing organisations, in time for them to stop the problem causing excessive disruption.
- **Loss reduction** This covers any action that reduces the catastrophic consequences of failure by removing the resources that are likely to suffer those consequences. For example, the road signs that indicate evacuation routes in the event of severe weather, or the fire drills that train employees in how to escape in the event of an emergency, may not reduce the consequences of failure on buildings or physical facilities, but can dramatically help in reducing loss of life or injury.
- **Substitution** This means compensating for failure by providing other resources that can substitute for those rendered less effective by the failure. It is a little like the concept of redundancy that was described earlier, but does not always imply excess resources if a failure has not occurred. For example, in a construction project, the risk of encountering unexpected geological problems may be mitigated by the existence of a separate work plan and that is invoked only if such problems are found. The resources may come from other parts of the construction project, which will in turn have plans to compensate for their loss.

Table 14.1 gives some examples of each type of failure mitigation actions for three failures: the theft of money from one of a company's bank accounts; the failure of a new product technology to work adequately during the new product development process; and the outbreak of fire at a business premises.

Table 14.1 Failure mitigation actions for three failures

	Type of failure		
Failure mitigation actions	***Financial failure – theft from company account***	***Development failure – new technology does not work***	***Emergency failure – fire at premises***
Mitigation planning	Identify different types of theft that have been reported and devise mitigation actions including software to identify anomalous account behaviour	Identify possible types of technology failure and identify contingency technologies, together with plans for accessing contingency technologies	Identify fire hazards and methods of detecting, limiting and extinguishing fires
Economic mitigation	Insure against theft and possibly use several different accounts	Invest in, or form partnership with, supplier of alternative technology	Insure against fire and have more, smaller, premises
Containment (spatial)	'Ring fence' accounts so a deficit in one account cannot be made good from another account	Develop alternative technological solutions for different parts of the development project so that failure in one part does not affect the whole project	Install localised sprinkler systems and fire door barriers
Containment (temporal)	Invest in software that detects signs of possible unusual account behaviour	Build-in project milestones that indicate the possibility of eventual development failure	Install alarm systems that indicate the occurrence of fire to everyone who may be affected (including in other premises)
Loss reduction	Build in transfer delays until approval for major withdrawals has been given, also institute plans for recovering stolen money	Ensure the development project can use old technology if new one does not work	Ensure means of egress and employee training are adequate
Substitution	Ensure that reserve funds and staff to manage the transfer can be speedily brought into play	Have fall-back work package for devoting extra resources to overcome the new technology failure	Ensure back-up team that can take-over from premises rendered inoperative by fire

EXAMPLE

Tiny and not so tiny recalls for General Motors[5]

According to the data analytics company Teradata Corp, car recalls cost the automotive industry somewhere between $45 and $50 billion per annum. Which is why automotive giants including General Motors, Nissan and Toyota have looked to build sophisticated data-mining operations so they can move towards what industry insiders call 'tiny recalls'. When faced with an emerging potential mechanical problem, car manufacturers have traditionally had two choices. The first option is to issue a general recall for the model in question. But there is a major financial and reputational cost to this, as well as significant inconvenience for customers, many of whom won't actually have a fault in their car. The second option is to ignore the fault and in doing so accept the risk that a very small number of customers may have problems (even serious accidents). While this may be justified in short-term financial (if not ethical) terms, the longer-term consequences could be disastrous; especially if the media exposes the fact. Now, there is a third option available.

The rise of bar-coding, radio-frequency identification (RFID), and more advanced information systems has allowed automakers to track components in far greater detail. The result is that faults can often be isolated not only to a model range, but to a specific set of components, made in particular periods of time, even by a particular shift at the factory. This then enables a far more focused recall of only the cars that are most likely to have a defect. For example, when a braking fault was identified on the Chevrolet Volt, General Motors (GM) was able to identify all 'at-risk' cars through its advanced parts tracking, and issue a 'tiny recall'. (In this case, just four cars in the whole of the US.) However, prevention is better than cure. Many auto manufacturers while seeking to target recalls, still face large recalls brought about by quality problems. For example, during 2014, GM announced the recall of a staggering 29 million cars and vans in North America, a number greater than its total US sales between 2005 and 2013. One of the largest recalls was for ignition switch defects, which included models as old as 1997.

DIAGNOSTIC QUESTION

Have failure recovery measures been implemented?

Not all failure can be prevented or mitigated. So operations need to plan carefully how they might recover if failure does occur and when negative impacts have been experienced. Yet even where the customer sees a failure, it may not necessarily lead to dissatisfaction; customers may even accept that things occasionally do go wrong. If there is a metre of snow on the train lines, or if the restaurant is particularly popular, we may accept that the product or service does not work. It is not necessarily the failure itself that leads to dissatisfaction but often the organisation's response to the failure. Mistakes may be inevitable, but dissatisfied customers are not. For example, when a flight is delayed by five hours, there is considerable potential for dissatisfaction. But if the airline informs passengers that the aircraft has been delayed by a cyclone at its previous destination and that arrangements have been made for accommodation at a local hotel with a complimentary meal, passengers might then feel that they have been well treated and even recommend that airline to others. As such, a good recovery can turn frustrated customers into loyal ones.

The complaint value chain

The complaint value chain helps us to visualise the potential value of good recovery at different stages (see Figure 14.11). In Figure 14.11(a) an operation provides service to 5,000 customers, but that 20 per cent experiences some form of failure. Of these 1,000 customers, 40 per cent decide not to complain, perhaps because it seems like more trouble than it's worth or because the complaint processes are too convoluted. Evidence suggests that around 80 per cent of these non-complainers will switch to an alternative service provider (of course the precise switching percentage will depend on the number of alternatives in the market and the ease of switching). Another group of the 1,000 customers who experienced a failure do decide to complain, in this case 60 per cent. Some will be satisfied (in this case, 75%) and others will not be (in this case, 25%). Dissatisfied complainers will generally leave the organisations

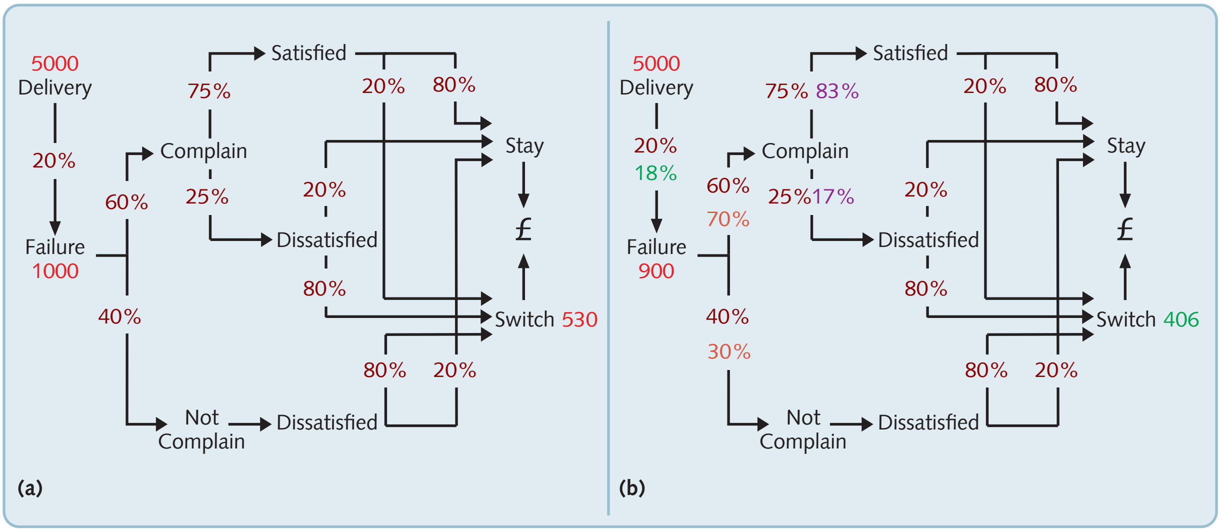

Figure 14.11 Complaint value chain (a) initial value chain and (b) with small improvements to each step

(for this example 80% to 20%) while satisfied complainers will tend to remain loyal (again, in this case 80% to 20%). So assuming these percentages are correct, for every 5,000 customers processed by this particular service operation, 530 will switch.

Now let's assume that the operations manager decides to invest in small improvement to all stages in the complaint value chain. In Figure 14.11(b) the company has reduced its failures from 20 per cent to 18 per cent (still very poor of course); and has encouraged more customers who experienced a failure to come forward and complain. So the percentage complaining has risen from 60% to 70%. It has also made sure that a higher proportion (in this case, from 75% to 83%) of those who do make the effort to complain are satisfied. The end result is that the number of lost customers falls from **530** to **406.** Assuming that an extra 124 customers have received value that is equal, or more than, the costs of improvements, the organisation is making a good investment in its recovery and prevention efforts. What is important to understand here is how a relatively small improvement across the failure and complaint process can have such a significant impact on customer loyalty and switching.

The recovery process

Recovery needs to be a planned process. Organisations therefore need to design appropriate responses to failure, linked to the cost and the inconvenience caused by the failure to their customers. Such recovery processes need to be carried out either by empowered front-line staff or by trained personnel who are available to deal with recovery in a way which does not interfere with day-to-day service activities. Figure 14.12 illustrates a typical recovery sequence.

Discover

The first thing any manager needs to do when faced with a failure is to discover its exact nature. Three important pieces of information are needed: first of all, what exactly has happened; second, who will be affected by the failure and, third, why did the failure occur? This last point is not intended to be a detailed inquest into the causes of failure (that comes later) but it is often necessary to know something of the causes of failure in case it is necessary to determine what action to take.

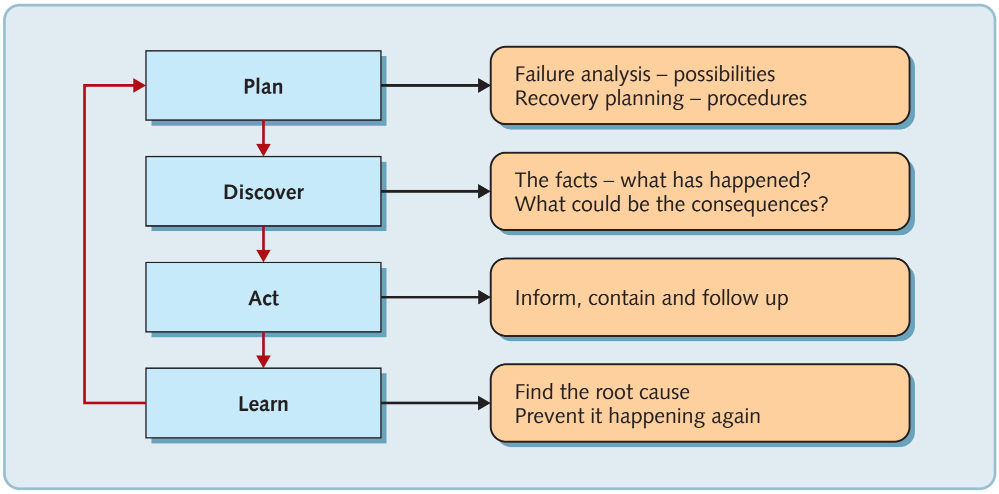

Figure 14.12 Recovery sequence for minimising the impact from failure

Act

The discover stage could only take minutes or even seconds, depending on the severity of the failure. If the failure is a severe one with important consequences, we need to move on to doing something about it quickly. This means carrying out three actions, the first two of which could be carried out in reverse order, depending on the urgency of the situation. First, tell the significant people involved what you are proposing to do about the failure. In service operations, this is especially important where the customers need to be kept informed. In all operations, however, it is important to communicate what action is going to happen so that everyone can set their own recovery plans in motion. Second, the effects of the failure need to be contained in order to stop the consequences spreading and causing further failures. The precise containment actions will depend on the nature of the failure. Third, there needs to be some kind of follow-up to make sure that the containment actions really have contained the failure.

Learn

As discussed earlier in this chapter, the benefits of failure in providing learning opportunities should not be underestimated. In failure planning, learning involves revisiting the failure to find out its root cause and then engineering out the causes of the failure so that it will not happen again.

Plan

Learning the lessons from a failure is not the end of the procedure. Operations managers need formally to incorporate the lessons into their future reactions to failures. This is often done by working through ‘in theory’ how they would react to failures in the future.

Critical commentary

The idea that failure can be detected through in-process inspection is increasingly seen as only partially true. Although inspecting for failures is an obvious first step in detecting them, it is not even close to being 100 per cent reliable. Accumulated evidence from research and practical examples consistently indicates that people, even when assisted by technology, are not good at detecting failure and errors. This applies even when special attention is being given to inspection. No one is advocating abandoning inspection as a failure detection mechanism. Rather it is seen as one of a range of methods of preventing failure. Much of the previous discussion surrounding the prevention of failure has assumed a 'rational' approach. In other words, it is assumed that operations managers and customers alike will put more effort into preventing failures that are either more likely to occur, or more serious in their consequences. Yet this assumption is based on a rational response to risk. In fact, being human, managers often respond to the perception of risk rather than its reality.

SUMMARY CHECKLIST

- ☐ Does the business have an operations and process resilience policy?
- ☐ Have all potential sources of failure been identified?
- ☐ Have any future changes in the sources of failure been identified?
- ☐ Is post-failure analysis carried out when failure does occur?
- ☐ Has the likelihood of each potential failure been assessed?
- ☐ Has the possibility of non-evident failures been addressed?
- ☐ Have the impact of all potential sources of failure been assessed?
- ☐ Are techniques such as failure mode and effect analysis (FMEA) used?
- ☐ Has due attention been paid to the possibility of designing out failure points?
- ☐ Has the idea of fail-safeing (poke yoke) been considered as a means of reducing the likelihood of failure?
- ☐ Have all approaches to process and technology maintenance been explored?
- ☐ Is the concept of redundancy economically viable for any potential failures?
- ☐ Does the operation have a failure mitigation plan?
- ☐ Have the whole range of mitigation actions been thoroughly evaluated?
- ☐ Are specific plans in place for the use of each type of mitigation action?
- ☐ Is a well-planned recovery procedure in place?
- ☐ Does the recovery procedure cover all the steps of discover, act, learn and plan?

CASE STUDY

Slagelse Industrial Services

Slagelse Industrial Services (SIS) had become one of Europe's most respected die caster of zinc, aluminium and magnesium parts supplier for hundreds of companies in many industries, especially automotive and defence. The company cast and engineered precision components by combining the most modern production technologies with precise tooling and craftsmanship. Slagelse Industrial Services (SIS) began life as a classic family firm by Erik Paulsen, who opened a small manufacturing and die-casting business in his hometown of Slagelse, a town in east Denmark, about 100 km southwest of Copenhagen. He had successfully leveraged his skills and passion for craftsmanship over many years, while serving a variety of different industrial and agricultural customers. His son, Anders had spent nearly ten years working as a production engineer for a large automotive parts supplier in the UK, but eventually returned to Slagelse to take-over the family firm. Exploiting his experience in mass manufacturing, Anders spent years building the firm into a larger scale industrial component manufacturer but retained his father's commitment to quality and customer service. After 20 years, he sold the firm to a UK-owned industrial conglomerate and within ten years it had doubled in size again and now employed in the region of 600 people and had a turnover approaching £200 million. Throughout this period the firm had continued to target their products into niche industrial markets where their emphasis upon product quality and dependability meant they were less vulnerable to price and cost pressures. However, in 2009, in the midst of difficult economic times and widespread industrial restructuring, they had been encouraged to bid for higher volume, lower margin work. This process was not very successful but eventually culminated in a tender for the design and production of a core metallic element of a child's toy (a 'transforming' robot).

Interestingly, the client firm, Alden Toys, was also a major customer for other businesses owned by SIS's corporate parent. They were adopting a preferred supplier policy and intended to have only one or two purchase points for specific elements in their global toy business. They had a high degree of trust in the parent organisation and on visiting the SIS site were impressed by the firm's depth of experience and commitment to quality. In 2010, they selected SIS to complete the design and begin trial production.

'Some of us were really excited by the prospect . . . but you have to be a little worried when volumes are much greater than anything you've done before. I guess the risk seemed okay because in the basic process steps, in the type of product if you like, we were making something that felt very similar to what we'd been doing for many years.' (SIS operations manager)

'Well obviously we didn't know anything about the toy market but then again we didn't really know all that much about the auto industry or the defence sector or any of our traditional customers before we started serving them. Our key competitive advantage, our capabilities, call it what you will, they are all about keeping the customer happy, about meeting and sometimes exceeding specification.' (SIS marketing director)

The designers had received an outline product specification from Alden Toys during the bid process and some further technical detail afterwards. Upon receipt of this final brief, a team of engineers and managers confirmed that the product could and would be manufactured using an up-scaled version of current production processes. The key operational challenge appeared to be accessing sufficient (but not too much) capacity. Fortunately, for a variety of reasons, the parent company was very supportive of the project and promised to underwrite any sensible capital expenditure plans. Although this opinion of the nature of the production challenge was widely accepted throughout the firm (and shared by Alden Toys and SIS's parent group) it was left to one specific senior engineer to actually sign both the final bid and technical completion documentation. By early 2011, the firm had begun a trial period of full volume production. Unfortunately, as would become clear later, during this design validation process SIS had effectively sanctioned a production method that would prove to be entirely inappropriate for the toy market, but it was not until 12 months later that any indication of problems began to emerge.

Throughout both North America and Europe, individual customers began to claim that their children had been 'poisoned' while playing with the end product. The threat of litigation was quickly leveled at Alden Toys and the whole issue rapidly became a 'full-blown' child health scare. A range of pressure groups and legal damage specialists supported and acted to aggregate the individual claims. Although similar accusations had been made before, the litigants and their supporters focused in on the recent changes made to the production process at SIS and in particular the role of Alden Toys in managing their suppliers.

'*. . . it's all very well claiming that you trust your suppliers but you simply cannot have the same level of control over another firm in another country. I am afraid that this all comes down to simple economics, that Alden Toys put its profits before children's health. Talk about trust . . . parents trusted this firm to look out for them and their families and have every right to be angry that boardroom greed was more important!*' (Legal spokesperson for US litigants when being interviewed on UK TV consumer rights show).

Under intense media pressure, Alden Toys rapidly convened a high-profile investigation into the source of the contamination. It quickly revealed that an 'unauthorised' chemical had been employed in an apparently trivial metal cleaning and preparation element of the SIS production process. Although when interviewed by the US media, the parent firm's legal director emphasised there was '*no causal link established or any admission of liability by either party*', Alden Toys immediately withdrew their order and began to signal an intent to bring legal action against SIS and its parent. This action brought an immediate end to production in this part of the operation and the inspection (and subsequent official and legal visits) had a crippling impact upon the productivity of the whole site. The competitive impact of the failure was extremely significant. After over a year of production, the new product accounted for more than a third (39%) of the factory's output. In addition to major cash-flow implications, the various investigations took up lots of managerial time and the reputation of the firm was seriously affected. As the site operations manager explained, even their traditional customers expressed concerns.

'*It's amazing but people we had been supplying for thirty or forty years were calling me up and asking '[Manager's name] what's going on?' and that they were worried about what all this might mean for them . . . these are completely different markets!*'

QUESTIONS

1 What operational risks did SIS face when deciding to become a strategic supplier for Alden Toys?

2 What control problems did they encounter in implementing this strategy (pre and post investigation)?

APPLYING THE PRINCIPLES

Some of these exercises can be answered by reading the chapter. Others will require some general knowledge of business activity and some might require an element of investigation. **Model answers for the first two questions are available on the student companion website.** For model answers to the other questions in this section, please ask your tutor/lecturer.

1 Revisit the six examples of failure described at the beginning of the chapter. Compile a table that summarises your view of:

(a) The reasons for the main failure

(b) How it might have been prevented

(c) The result (damage) of the failure

2 Wyco is a leading international retailer selling clothing and accessories with stores throughout the United States, Europe and the Middle East. The countries from which it sources its products include Sri Lanka, Bangladesh, India and Vietnam. It was shocked when a British newspaper reported that an unauthorised subcontractor had used child workers to make some of its products at a factory in Delhi. In response, Wyco immediately issued a statement.

'*Earlier this week, the company was informed about an allegation of child labor at a facility in India that was working on one product for Wyco. An investigation was immediately launched. The company noted that a very small portion of a particular order placed with one of its vendors was apparently subcontracted to an unauthorised subcontractor without the company's knowledge or approval. This is in direct violation of the company's agreement with the vendor under its Code of Vendor Conduct*.' The company CEO said, '*We strictly prohibit the use of child labor. This is a non-negotiable for us – and we are deeply concerned and upset by this allegation. As we've demonstrated in the past, Wyco has a history of addressing challenges like this head-on. Wyco ceased business with 20 factories due to code violations. We have 90 people located around the world whose job is to ensure compliance with our Code of Vendor Conduct. As soon as we were alerted to this situation, we stopped the work order and prevented the product from being sold in stores. While violations of our strict prohibition on child labor in factories that produce product for the company are extremely rare, we have called an urgent meeting with our suppliers in the region to reinforce our policies. Wyco has one of the industry's most comprehensive programs in place to fight for workers' rights. We will continue to work with stakeholder organisations in an effort to end the use of child labor.*'

(a) 'Being an ethical company isn't enough anymore. These days, leading brands are judged by the company they keep.' What does this mean for Wyco?

(b) When Wyco found itself with this supply chain problem, did it respond in the right way?

3 One cause of aircraft accident is 'controlled flight into ground'. Predominantly, the reason for this is not mechanical failure but human failure such as pilot fatigue. Boeing, which dominates the commercial airline business, has calculated that over 60 per cent of all the accidents which have occurred in the past 10 years had flight crew behaviour as their 'dominant cause'. For this type of failure to occur, a whole chain of minor failures must happen. First, the pilot at the controls has to be flying at the wrong altitude – there is only one chance in a thousand of this. Second, the co-pilot would have to fail to cross-check the altitude – only one chance in a hundred of this. The air traffic controllers would have to miss the fact that the plane was at the wrong altitude (which is not strictly part of their job) – a one-in-ten chance. Finally, the pilot would have to ignore the ground proximity warning alarm in the aircraft (which can be prone to give false alarms) – a one-in-two chance.

(a) What are your views on the quoted probabilities of each failure described above occurring?

(b) How would you try to prevent these failures occurring?

(c) If the probability of each failure occurring could be reduced by a half, what would be the effect on the likelihood of this type of crash occurring?

4 (True story) The light bulb in the men's lavatories of a factory finally burnt out after 70 years of operation. The manager at the firm said, 'It is actually a little bit sad. I joined the firm when I was fifteen (he is now sixty three) and it was there then.' In fact, the bulb had survived bombs dropped in the Second World War that had devastated neighbouring buildings, the army firing its guns in the next-door park and punk band 'the Clash' playing at a neighbouring venue, which caused residents to complain that their windows were being shaken by the noise. When it did eventually fail, the firm had it mounted on a stand and gave it a place of honour. More remarkable, they tracked down the original supplier of the bulb. He had also survived. He was 99 years old and still remembered selling them. Does this incident invalidate the use of failure data in estimating component life?

5 An automated sandwich-making machine in a food manufacturer's factory has six major components, with individual reliabilities as shown in Table 14.2

(a) What is the reliability of the whole system?

(b) If it is decided that the wrapper in the automated sandwich-making machine is too unreliable and a second wrapper is needed, which will come into action if the first one fails, what will happen to the reliability of the machine?

Table 14.2 Individual reliabilities of major components

Component	*Reliability*
Bread slicer	0.97
Butter applicator	0.96
Salad filler	0.94
Meat filler	0.92
Top slice of bread applicator	0.96
Wrapper	0.91

Notes on chapter

1 Sources include: Zetter, K. (2016) 'That insane, $81M Bangladesh Bank heist? Here's what we know', *Wired*, 17 June; Caroll, C. (2009) 'Defying a reputational crisis – Cadbury's salmonella scare: why are customers willing to forgive and forget?, *Corporate Reputation Review*, vol.12, Issue 1, 64–82; *Raconteur* (2013), 'Business risk strategies', 7 May; *BBC News*, 2013 'Findus beef lasagne contained up to 100% horsemeat, FSA says', 7 February; Szu Ping Chan, and agencies (2013) 'Timeline: how G4S's bungled Olympics security contract unfolded', *The Telegraph*, 21 May; Neate, R. (2013) 'G4S profits tumble on Olympic failings', The *Guardian*, 13 March; Tovey, A. (2017) 'VW attacked by MPs over failure to release findings of 'dieselgate' investigation', *The Telegraph*, 22 March; Woodman, P. (2008) 'Disastrous opening day for Terminal 5', *The Independent*, 27 March.

2 Source: The Economist (1994) 'Air crashes, but surely . . . ', 4 June.

3 Source: Graham-Harrison, E. (2015) 'M&S and others supplied by factories that mistreat workers, rights group says', The *Guardian*, 12 March.

4 Sources include; Blastland, M. and Spiegehalter, D. (2013) *The Norm Chronicles; Stories and Numbers About Danger*, Profile; Spiegelhalter, D. (2014) 'The power of the MicroMort', *British Journal of Obstetrics and Gynaecology*, vol. 121, 662-663; *The Economist* (2013), 'Making sense of the statistics that riddle our days', 22 June.

5 Sources include; Nelson, G. (2013) 'How data mining helped GM limit a recall to four cars', *Automotive News*, 28 October 28; Valdes-Dapena, P. (2014) 'GM: Steps to a recall nightmare', *CNN News*, Retrieved 20 October; Pollock, L. (2014) 'What you need to know about the GM recalls', *Wall Street Journal*, 2 April; Bennett, J. (2014) 'GM to recall 8.45 million more vehicles in North America', *Wall Street Journal*, 30 June.

TAKING IT FURTHER

Hopkin, P. (2017) ***Fundamentals of Risk Management: Understanding, evaluating and implementing effective risk management (4th edn), Kogan Page.*** *A comprehensive introduction to risk with good coverage of many core frameworks.*

Hubbard, D.W. (2009) ***The Failure of Risk Management: Why it's Broken and How to Fix It, John Wiley & Sons.*** *An interesting read, particularly for those who like the critical commentaries in this book! A polemic, but one that is clearly written.*

Smith, D.J. (2011) ***Reliability, Maintainability and Risk, Butterworth-Heinemann.*** *A comprehensive and excellent guide to all aspects of maintenance and reliability. The book has a good mix of qualitative and quantitative perspective on the subject.*

Waters, D. (2012) ***Supply Chain Risk Management: Vulnerability and Resilience in Logistics (2nd edn), Chartered Institute of Logistics and Transportation.*** *Provides a very detailed and practical guide to considering risks within operations and supply chains.*

15

Project management

Introduction

In this chapter, we examine the management of projects. Projects come in all shapes and sizes, with differences in scale, uncertainty, complexity, novelty, technology and pace. Yet, in many respects, projects share key characteristics that make their management tasks broadly universal. First, managers must understand the fundamental characteristics of a project and the likely implications of these characteristics for management. Second, they must be able to identify and manage key project stakeholders. Third, they must be able to define, plan and control projects through their lifecycle, while balancing competing performance objectives and competing (internal and external) stakeholder requirements. Figure 15.1 shows where project management fits in the overall structure of this book.

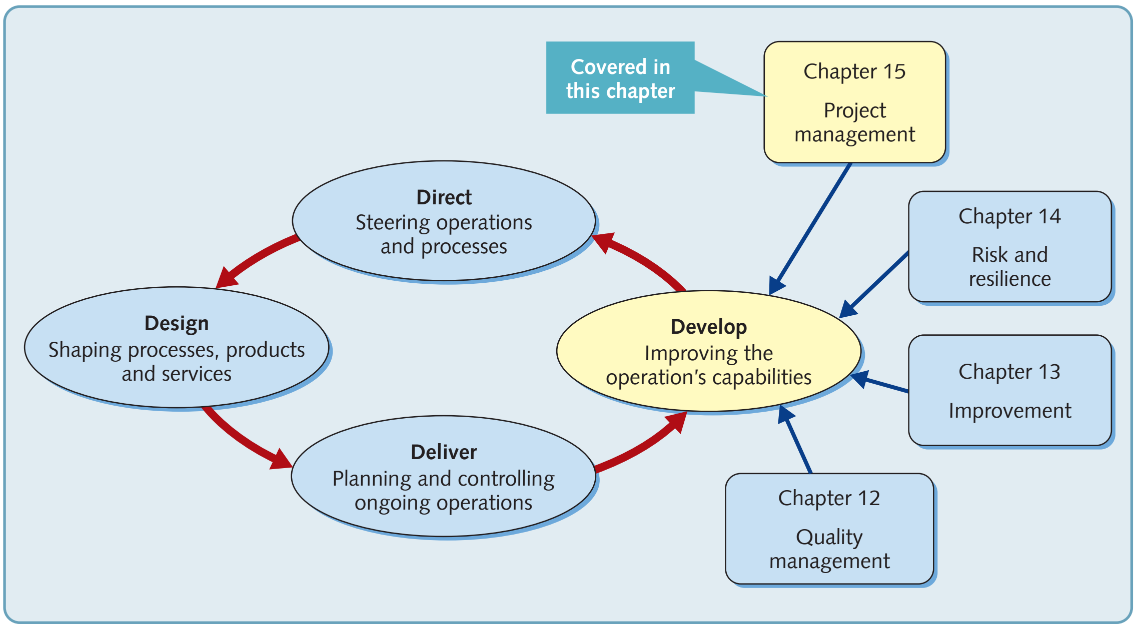

Figure 15.1 **Project management is the activity of defining, planning and controlling projects**

EXECUTIVE SUMMARY

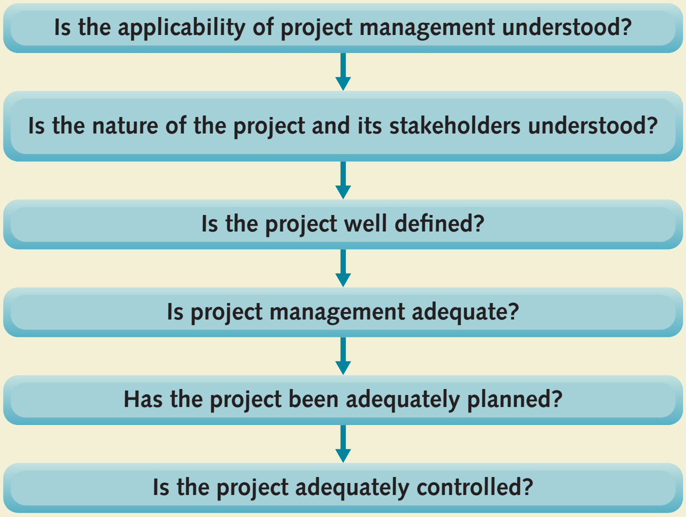

Is the applicability of project management understood?

A project is a set of activities dedicated to achieving a specific goal within a set time-frame, to certain specifications, using a defined group of resources. As such, projects are highly customised or unique, and often face high levels of risk and uncertainty. Project management is not just the province of project managers. It is a ubiquitous task of operations management. People in all types of operation will get involved in some element of project management because it is the activity of defining, planning, controlling and learning from projects. It is also concerned with the operations task of effectively balancing the so-called 'iron triangle' of quality/deliverables, time and cost objectives. Finally, from an organisational perspective, project management involves coordinating these lifecycles and objectives across multiple functions.

Is the nature of the project and its stakeholders understood?

The process of managing a project begins by understanding its fundamental characteristics. In this chapter, we explore three alternative methods that managers can use to understand the nature of a project: differentiating by volume and variety characteristics; differentiating by scale, complexity and uncertainty characteristics; and differentiating using the 'diamond model' of novelty, technology, complexity and pace. Once project characteristics are understood, project managers must identify individuals, groups, or entities that have an interest in the project process or outcome. They must then decide how to engage with different stakeholders and how best to manage competing needs.

Is the project well defined?

Defining the project involves three related activities: defining project objectives, project scope and project strategy. Most projects can be defined by the relative importance of three objectives. These are: cost – keeping the overall project to its original budget; time – finishing the project by the scheduled finish time; and quality – ensuring that the project outcome is as was originally specified. The project scope identifies its work content and its outcomes. The task is not only concerned with what the project will do, but also what it will not do, that is, its limits and exclusions. The project strategy describes the general way in which the project is going to meet its objectives, including significant project milestones and 'stagegates'.

Is project management adequate?

Because of their complexity and the involvement of many different parties, projects need particularly careful management. In fact, project management is seen as an especially demanding role with a very diverse set of skills including technical project management knowledge, interpersonal skills and leadership ability. Very often, project managers need the ability to motivate staff who not only report into managers other than themselves, but also divide their time between several different projects.

Has the project been adequately planned?

Project planning involves determining the cost and duration of the project and the level of resources that it will need. In more detail, it involves identifying the start and finish times of individual activities within the project. Generally, the five stages of project planning include identifying activities, estimating times and resources, identifying relationships and dependencies between activities, identifying time and resource schedule constraints and fixing the final schedule. Critical path analysis (CPA) and Program evaluation and review technique (PERT) are often used to aid the project planning process.

Is the project adequately controlled?

Project control involves monitoring the project in order to check its progress, assessing the performance of the project against the project plan and, if necessary, intervening in order to bring the project back to plan. The process often involves continually assessing the progress of the project in terms of budgeted expenditure and progress towards meeting the project's final goal. It may also involve deciding when to devote extra resources to accelerating (also known as crashing) individual activities within the project.

DIAGNOSTIC QUESTION

Is the applicability of project management understood?

Projects in operations management come in all shapes and sizes. Some are associated with the development of specific products or services. Some are capital investment or infrastructure projects, such as organising the installation of new buildings or process technologies. Some involve specific pre-planned organisational change that seeks to make a change to methods of working or reporting responsibilities. Some have a clear objective but a less well-defined plan for achieving the objective, such as some improvement initiatives. Some are large and complex, with multiple stakeholders, such as major public infrastructure initiatives. Others are short, small-scale and limited, such as a simple inventory check.

But, while projects vary a great deal, it is important to stress that project management is not just the province of project managers. It is a ubiquitous task of operations management. People in all types of operation will get involved in some element of project management. And, while projects vary a great deal, they do share a number of common features. All projects are mission focused – that is, they are dedicated to achieving a specific goal that should be delivered within a set timeframe, to certain specifications, using a defined group of resources. The result is that project outcomes are either unique or at least highly customised, involve many non-routine and complex tasks and therefore face relatively high levels of risk and uncertainty when contrasted with day-to-day operations. And it is these features that perhaps explain why so many projects fail in some way, with changed specifications (quality), severe delays (time), cost escalation (cost) and major disputes between key stakeholders commonplace. However, it is not only the innate complexity of projects that leads to many failing, it is often the lack of effective project management.

What is project management?

Project management is the activity of defining, planning and controlling, and learning from projects of any type. Going beyond this lifecycle perspective, project management is also concerned with effectively balancing quality/deliverables, time and cost objectives within the so-called 'iron triangle' (of quality, time and cost). Finally, from an organisational perspective, project management involves managing these lifecycles and performance objectives across multiple functions within an organisation. The activity of project management is very broad in so much as it could encompass almost all the operations and process management tasks described in this book. Partly because of this, it could have been treated almost anywhere within its direct, design, delivery, develop, structure. We have chosen to place it in the context of operations and process development because the majority of projects that managers will be engaged in are essentially improvement projects. Of course, many projects are vast enterprises with very high levels of resourcing, complexity and uncertainty that will extend over many years. Look around at the civil engineering, social, political and environmental success (and failures) to see the evidence of major projects. Such projects require professional project management involving high-level technical expertise and management skills. But, so do the smaller, yet important, projects that implement the many and continuous improvements that will determine the strategic impact of operations development. This is why it is equally important to take a rigorous and systematic approach to managing projects regardless of their size and type.

'Projects' and 'programmes'

It is also worth pointing out the distinction between 'projects' and 'programmes'. A programme, such as a *continuous* improvement programme, has no defined end point. Rather it is an ongoing process of change. Individual projects, such as the development of a new training workshop, may be individual sub-sections of the overall programme. Programme management will overlay and integrate the individual projects. Generally, it is a more difficult task in the sense that it requires resource coordination, particularly when multiple projects share common resources, as emphasised in the following quotation. *'Managing projects is, it is said, like juggling three balls – cost, quality, and time. Programme management. . . is like organising a troupe of jugglers all juggling three balls and swapping balls from time to time.'*[1]

The following two projects illustrate some of the issues in project management.

EXAMPLE

Imagineering projects at Disney[2]

Ever since the creation of famous characters such as Mickey Mouse and Snow White, the Disney Corporation has been synonymous with innovative quality entertainment. At The Walt Disney Studios and Parks, persuasive stories come alive and technology and artistry come together to delight visitors and audiences alike. Nowhere is this truer than Walt Disney's 'Imagineering' operation ('Imagineering' is a word that combines 'imagination' and 'engineering'), which, in effect, is the research and development arm of Walt Disney Parks and Resorts. It is responsible for creating many of the more significant attractions, shows, fireworks displays and parades at the Disney theme parks.

Their complex and often innovative attractions will be closely scrutinised by thousands or millions of Disney's 'guests', so it is vital to pay attention to details. A famous Walt Disney saying is that a guest may not notice a specific (sometimes tiny) detail, but he or she will notice when the detail isn't there. This means that they need to be created with skill, creativity and, equally importantly, professional project management. Walt Disney World's Imagineering Department's project managers work with the interactive technologies, the park staff, special-effects wizards, digital designers and others to create an interactive experience for the guests. Although many of the projects are technical in nature, project managers work with a wide range of different disciplines from construction to marketing.

David Van Wyk is the vice-president of project management for Walt Disney Imagineering and he fully understands the importance of effective project management. *'Without it'*, he says, *'how can we be as relevant tomorrow as we are today? How can we meet and exceed guest expectations in a changing world? We have somewhere between 140 and 150 different skill sets in Imagineering, including engineers, creative staff, artists, architects, accountants, writers, theme and new media specialists, and more. A culture of interdisciplinary coordination with diverse stakeholders aims to interact and socialise to understand issues and problems. We have developed and keep working on a culture of collaboration. We do not have an NIH (not invented here) culture.'* The Imagineering group is very much aware of the implications of the MacLeamy Curve (see the text) that highlights the idea that cost increases while the ability to change decreases over the life of the project. So it is important to solve issues earlier in the design process, when it's more economical to make changes, especially those involving equipment. Which is why, says David Van Wyk, '*that we seek to incorporate more peer review earlier* *in the engineering-design process. We also look for on-time delivery, getting it right before it gets to the field, with a strong start, strong finish, and careful resource allocation.'* All of which makes it important to develop excellent relationships with partners and make sure that they, and all stakeholders, fully subscribe to the objectives of predictability, collaboration, impeccable coordination, prompt decision-making, collective quality and just-in-time delivery.

EXAMPLE

Crossrail, Europe's biggest infrastructure project[3]

It was an idea that took over a hundred years to be approved and 10 years to complete – not surprisingly, the £14.8 billion Crossrail project was one of the most complex and ambitious that London had known, and was, at the time, Europe's biggest infrastructure project. The new railway is forecast to carry 200 million passengers a year, providing a 10 per cent increase in central London's rail capacity, providing more frequent and reliable train journeys for London's growing population. It will serve 40 stations, including eight new subsurface stations and two new above-ground stations, with 42 km of new tunnels, and using 200-metre-long new trains. Crossrail's eight tunnel-boring machines, each weighing 1,000 tonnes, spent 3 years burrowing under London to construct the new tunnels. It was a huge technical challenge. To construct the tunnels, which were completed in 2015, massive tunnel-boring machines worked around the clock, moving through ground already honeycombed with networks of sewer lines, water and gas mains, building foundations and London Underground tunnels dating to the 1860s. But the idea was not new. It was during the 1880s when the UK Government gave permission for the Regents Canal & Railway Company to create a surface line between west London and the busy dockland area in East London. The plan was discarded, but a similar idea re-emerged in 1944, only to be abandoned again. Through the 1990s the project was the subject of several policy changes and cost–benefit discussions, and it was not until 2008 that the project was finally approved, with construction starting in 2009. The governance of the project was always seen as important. Responsibility for delivering the final project was given to Crossrail Ltd, an independent company with all of its shares owned by Transport for London (TfL). They reported to a 'sponsor group' with representatives of TfL and Government departments. The sponsor group was responsible for defining the scope of the project while holding the budget for the project. Crossrail Ltd had to work within the budget and manage the details of the project including scheduling and control activities, cost control and making sure the reporting regime was a true reflection of what was really happening. A succession of 'gateways' were also agreed. These were deadlines by which Crossrail Ltd was required to demonstrate it had met stages of completion. As the project progressed and gateways were passed, Crossrail was given more power to award large contracts itself. Some of the project's managers attribute its success to the tight management of both costs and schedules, as well as the ability to adapt; for example, efficiencies of more than £1 billion were made within the programme budget by extending the delivery time for a major part of the project (the central tunnel works) by around a year. However, others are more critical. Professor Gil of the Alliance Manchester Business School argues that the project's governance structure failed to prevent the spending of £5 billion of contingency funding that had been included in the budget (later reduced to £4 billion). '*The spending of contingencies becomes a self-fulfilling prophecy that politicians and managers can't prevent. If there was less money for a project available at the beginning, it might focus the stakeholders*' minds better.' Some have suggested that the lesson from the decision to separate the sponsor group from the delivery company, is that for such large and complex infrastructure projects, civil servants should be kept well away from delivery. '*But that is not totally true*', said James Stewart, chairman of global infrastructure at consultancy KPMG. '*The competencies required within Whitehall are not to run the day-to-day commercial procurements and project management*, *it is true. But there is a lot of skill in being a good client. You are monitoring the performance of the delivery body and not relying on their assurances that everything is fine. You are also managing the politics and policy environment*.'

What do they have in common?

In some respect these examples illustrate two types of projects. One organisation (Disney) is set in the private sector, projects are of relatively short duration (theme park rides need to be current), there will be several projects running concurrently, and although each project might be different, they probably have similar characteristics. The other (Crossrail) is a single public sector project, is incredibly complex and expensive, with multiple stakeholders, has been discussed for over a hundred years and has a duration of ten years. Yet in both types of projects many similar project management issues are evident. Both have an objective, a definable end result. Both are temporary in that they have a defined beginning and need a temporary concentration of resources that will be redeployed once their contribution has been completed. Both need to motivate the people involved in the project; both need a management (governance) structure to oversee the planning and controlling of the project(s). Finally, because no single project exists in isolation, both projects must manage a variety of stakeholders. For Disney, these may include the ride designers, suppliers, employees, shareholders, community groups, retailers, sponsors, existing and potential customers, safety regulators, and so on. For Crossrail, they will include numerous governmental bodies, private contractors, existing transport operations, their staff and the travelling public. As such, each project's objectives and scope need to be clarified with stakeholders, while detailed planning must then be carried out by the project manager and their project team. They are the ones who determine the time and resource commitments that the project will need, as well as identifying the long list of things that could go wrong.

DIAGNOSTIC QUESTION

Is the nature of the project and its stakeholders understood?

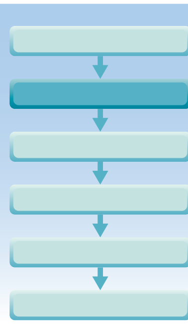

The management of any project should begin by understanding the nature of the project and its broad characteristics. It should then identify and manage the 'stakeholders' (also sometimes referred to as agents) who may affect, or be affected by, the success or failure of the project and its activities. Both of these elements of the project's environment will fundamentally impact how project managers go about ensuring that the project reaches a successful conclusion.

EXAMPLE

Halting the growth of malaria[4]

In 2000 the World Health Organization (WHO) set a challenging objective – to halt the growth of malaria. At the time, there were an estimated 300–500 million cases of malaria each year, with between 1.1 and 2.7 million deaths, with the largest proportion of these children under 5 years old. The WHO faced a hugely complex project management climate, with major political, economic, climatic and cultural impediments to success. And yet, by 2013 the reported number of malaria cases was down to 198 million and the number of deaths down to 580,000. At the heart of its success was a clear overriding vision that gained buy-in from the project's diverse set of stakeholders. Building on this, the WHO spent significant time understanding the complex project environment – the internal and external factors that might influence the success or failure of its various malaria projects worldwide. They also committed significant resources to objective setting, scoping and planning of their projects. Finally, in deploying and executing the various existing malaria-related technologies, both those focused on preventing malaria and those focused on curing people infected, the WHO and its partners relied heavily on careful project monitoring, milestones, and continuing stakeholder engagement to ensure that they were on track. The fight against malaria is far from over, but at least this preventable and curable disease is in decline.

Differentiating between projects

So far, we have shown what projects have in *common* – temporary activities, with specific and highly customised goals, within time, cost and quality requirements, usually involving many non-routine and complex tasks. However, it is also critical to understand *differences* between projects. To some extent we have already started to do this when we identified the differences between the two examples of projects earlier in the chapter – Disney and Crossrail. But it also requires an understanding of the other factors that will influence how the project is managed. We illustrate three complementary methods of distinguishing between different types of project:

> OPERATIONS PRINCIPLE
>
> ***The difficulty of managing a project is a function of (a) its volume and variety characteristics; (b) its scale, complexity, and uncertainty characteristics; and (c) its novelty, technology, complexity and pace characteristics.***

- the volume and variety characteristics of a project
- the scale, complexity and degree of uncertainty in the project
- the novelty, complexity, nature of technology (if any) and 'pace' of the project.

We will deal with these methods in turn.

Differentiating projects by their volume and variety characteristics

At a simple level, we can use the 'product-process' matrix, already explored in Chapter 5 of this book, to distinguish between projects based on their volume and variety characteristics. This is shown in Figure 15.2. Of course, all project processes are, by definition, in the top left corner of the matrix. But within that end of the 'natural diagonal' projects do vary. At the very top left-hand part of the matrix are projects that are genuinely '*first timers*' with a very high degree of uniqueness, a volume of one and infinite variety. With less uniqueness, higher volume and less variety, '*as before, but. . .* ' projects may share some of the attributes of previous projects, but may have new features where project managers have little or no previous experience to help guide them. With higher volume (therefore a greater degree of repetition) and lower variety, so-called '*paint by numbers*' projects, are relatively routine and predictable, and therefore (generally) more straightforward to manage.

Differentiating by projects by their scale, complexity and uncertainty characteristics

An alternative approach to distinguishing between projects is by considering their scale, complexity and uncertainty. This is shown in Figure 15.3. For example, a wedding planning project has (relatively) low levels of scale, complexity and uncertainty. The effect is that the management challenges of such a project are significantly different to developing the Airbus A380, which exhibited much higher levels of all three dimensions. The scale, complexity and uncertainty of such 'ground-breaking' projects demand far more sophisticated planning, greater and more flexible resources, and careful control if they are to be successful.

Differentiating using novelty, technology, complexity and pace

Yet another alternative (and very useful) way to distinguish between projects is to consider their relative novelty, technology, complexity and pace. Using the scales and terminology developed by Aaron Shenhar and Dov Dvir,[5] Figure 15.4 illustrates the profile of two projects – the development of the Airbus A380 and the WHO's malaria project described in the earlier example.

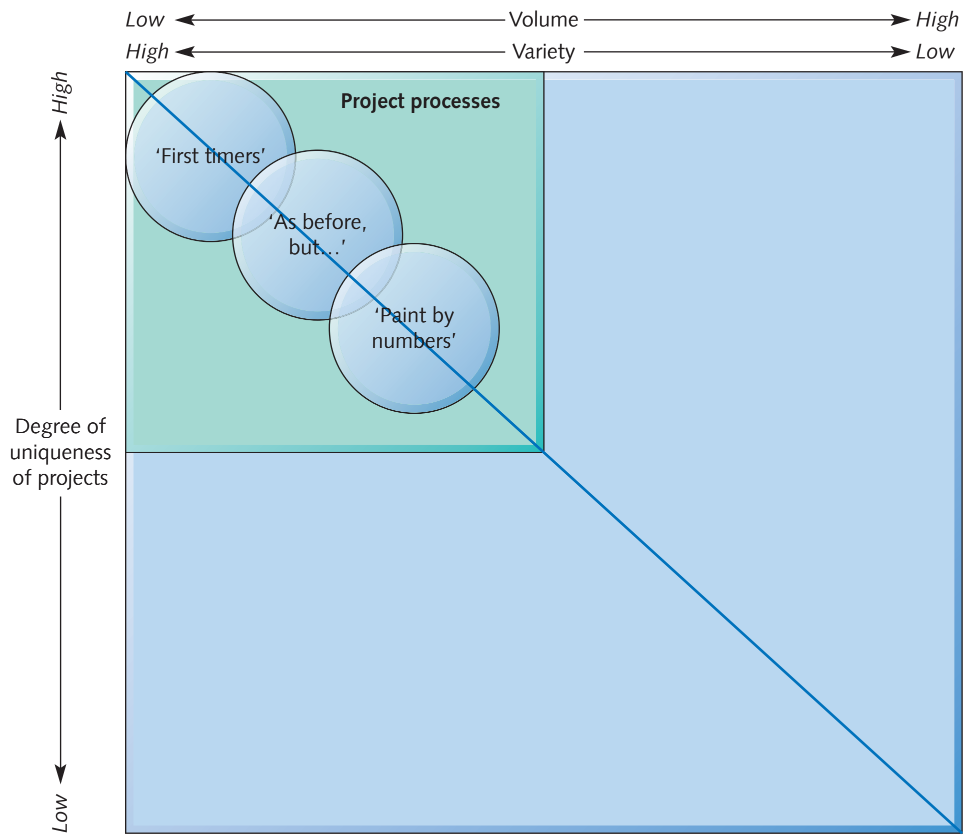

Figure 15.2 Differentiating projects using their volume and variety characteristics

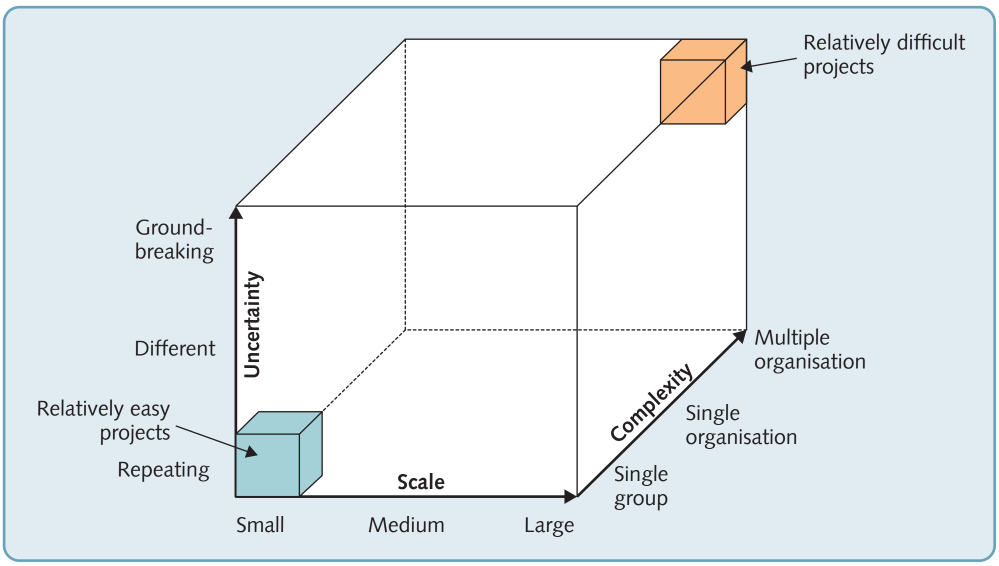

Figure 15.3 Differentiating projects by scale, complexity and uncertainty

The novelty dimension

This is concerned with how new the outcome of the project is to the customers or users (i.e. the market). On this scale, a derivative project is one that extends or improves an existing product, service or process; for example, developing a new version of a phone app. A platform project is

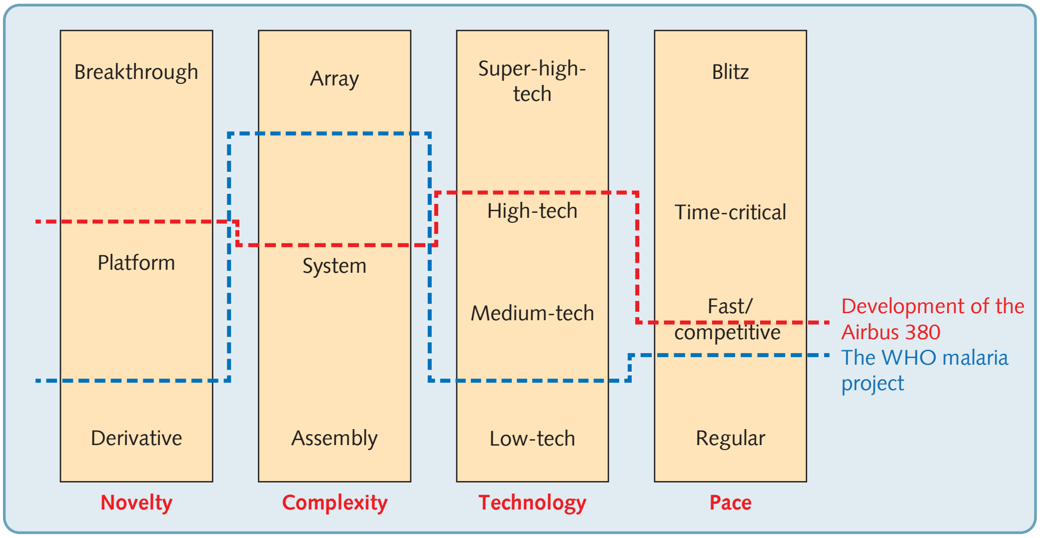

Figure 15.4 **Differentiating projects using novelty, technology, complexity and pace**

one that develops or produces new generations of existing products, services or processes to serve existing markets; for example, building a new generation of car. A breakthrough project is one that introduces a totally new product or service, uses a new idea, or uses a product or service in a way that customers have never seen before; for example, introducing the very first iPad.

The technology dimension

This is concerned with how much new technology is being used within the project. Low-technology projects have almost no new technology integrated, so designs can be 'frozen' (i.e. fixed) early on in the project. Medium-technology projects typically involve the integration of a single new technology, for example, the improvement of an existing product. This allows early design-freezes, but some testing, evaluation and corrections may be required as the project progresses. High-technology projects involve the integration of several new technologies and therefore must be flexible for a longer period of time to allow for integration and optimisation. Finally, super high-technology projects involve the integration of several (currently) non-existing technologies. This takes extended periods to develop and prove new technologies, so design-freeze typically occurs very late in the project.

The complexity dimension

This is concerned with how complex the system and its sub-systems are. At the lowest level, are 'assembly' projects, with self-contained 'components', or sub-projects that perform a function within a larger system. Examples may include developing a next-generation smartphone, creating a new undergraduate operations and process management module, or putting on a new play in a theatre. 'System' projects involve a collection of interactive elements and sub-elements. Examples include developing a new aircraft, constructing a new research and development facility, or developing a new portfolio of post-experience education within a university. While the sub-elements of the project have a common goal, the added complexity creates significantly higher coordination and integration problems. Finally, 'array' projects are 'systems of systems' – with each system having an independent function, but each with a common goal. Heathrow's Terminal 5 project, with its 16 major projects and 147 sub-projects, is a good example of an array project. Another is the South-to-North water diversion (南水北□□程) project in China, a multi-decade infrastructural mega-project expected to be completed in 2050 at a cost of $62 billion.

The pace dimension

This is concerned with how critical the time frame of the project is. Pace is not simply about speed – some projects have urgency but last for many years, others are not urgent but last a few weeks. For example, in May 1961, President John F. Kennedy delivered a speech to the US Congress in which he stated, *'I believe that this nation should commit itself to achieving the goal, before this decade is out, of landing a man on the moon and returning him safely to the earth.'* In doing so, he set a time frame that was to be critical to the ambitions of the moon-landing project. Some projects are regular pace in that they are not time critical. Many public works and internal projects fall into this category. Others are fast/competitive in that the completion on time is important for competitive advantage and leadership. Many business-related projects, such as new product/service introductions and new capacity development in the face of market growth, fall into this category. Time-critical projects have a specific window of opportunity and delays mean project failure. The space launch or the Olympic games are examples of this kind of project. Finally, 'blitz' projects have the utmost urgency and often occur through crises such as war, response to natural disasters, and fast response to business surprises.

Understanding and managing stakeholders

Once managers have understood the fundamental characteristics of a project, they must consider the stakeholders or agents who are likely to interact with the project, and who could play a critical role in its success or failure. Project stakeholders are those individuals, groups, or entities that have an interest in the project process or outcome. In other words, they affect, or are affected by, the project. Internal stakeholders include the client, the project sponsor, the project team, functional managers, contractors, and project support. External stakeholders (i.e. those outside of the core project, rather than outside of the organisation) include end users, suppliers, competitors, lobby groups, shareholders, government agencies and employees.

All projects will have stakeholders; complex projects will have many. They are likely to have different views on a project's objectives that may conflict with other stakeholders. At the very least, different stakeholders are likely to stress different aspects of a project. So, as well as an ethical imperative to include as many people as possible in a project from an early stage, it is often useful in preventing objections and problems later in the project. Moreover, there can be significant direct benefits from using a stakeholder-based approach. Project managers can use the opinions of powerful stakeholders to shape the project at an early stage. This makes it more likely that they will support the project, and can also improve its quality. Communicating with stakeholders early and frequently can ensure that they fully understand the project and understand potential benefits. Stakeholder support may even help to win more resources, making it more likely that projects will be successful. Perhaps most important, one can anticipate stakeholder reaction to various aspects of the project, and plan the actions that could prevent opposition, or build support.

OPERATIONS PRINCIPLE
All projects have stakeholders with different interests and priorities.

Managing stakeholders can be a subtle and delicate task, requiring significant social and, sometimes, political skills. But it is based on three basic activities: identifying stakeholders; determining the nature of different stakeholders; and prioritising and managing stakeholders.

Identifying stakeholders

Think of all the individuals, groups, or entities who affect or are affected by your work, who have influence or power over it, or have an interest in its successful or unsuccessful conclusion. Although stakeholders are not just individuals, ultimately you must communicate with people, so look to identify key individuals within a stakeholder organisation. In addition, even if one decides not to attempt to manage every identified stakeholder, the process of stakeholder mapping is still useful because it gets those working on a project to see the variety of competing forces at play in many complex projects.

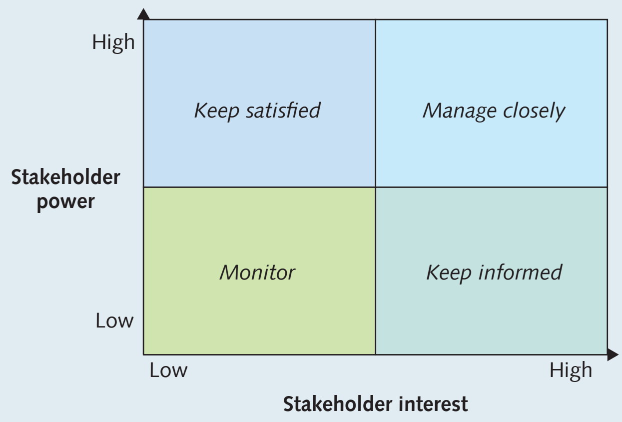

Figure 15.5 **The stakeholder power–interest grid**

Determining the nature of different stakeholders

Once all stakeholders have been identified, it is important to understand how they are likely to feel about and react to the project, and how best to communicate with them. One approach to discriminating between different stakeholders, and more important, how they should be managed, is to distinguish between their power to influence the project and their interest in doing so. Stakeholders who have the power to exercise a major influence over the project should never be ignored. At the very least, the nature of their interest, and their motivation, should be well understood. But not all stakeholders who have the power to exercise influence over a project will be interested in doing so, and not everyone who is interested in the project has the power to influence it. The power–interest grid, shown in Figure 15.5, classifies stakeholders simply in terms of these two dimensions. Although there will be graduations between them, the two dimensions are useful in providing an indication of how stakeholders can be managed in terms of four categories.

High-power, interested groups must be fully engaged, with the greatest efforts made to satisfy them. High-power, less interested groups require enough effort to keep them satisfied, but not so much that they become bored or irritated with the message. Low-power, interested groups need to be kept adequately informed, with checks to ensure that no major issues are arising. These groups may be very helpful with the detail of the project. Low-power, less interested groups need monitoring though without excessive communication. Some key questions that can help to understand high priority stakeholders include the following:

OPERATIONS PRINCIPLE

Different stakeholder groups will need managing differently.

- What financial or emotional interest do they have in the outcome of the project? Is it positive or negative?
- What motivates them most of all?
- What information do they need?
- What is the best way of communicating with them?
- What is their current opinion of the project?

- Who influences their opinions? Do some of these influencers therefore become important stakeholders in their own right?
- If they are not likely to be positive, what will win them around to support the project?
- If you don't think you will be able to win them.

DIAGNOSTIC QUESTION

Is the project well defined?

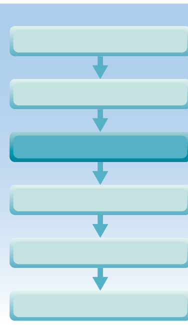

Before starting the complex task of planning and executing a project, it is necessary to be clear about exactly what the project is – its definition. This is not always straightforward, especially in projects with many stakeholders. Three different elements define a project:

- objectives: the end state that project management is trying to achieve
- scope: the exact range of the responsibilities taken on by project management
- strategy: how project management is going to meet its objectives.

Project objectives

Objectives help to provide a definition of the end point that can be used to monitor progress and identify when success has been achieved. They can be judged in terms of the five performance objectives: quality, speed, dependability, flexibility and cost. However, flexibility is regarded as a 'given' in most projects that, by definition, are to some extent one-offs, and speed and dependability are compressed to one composite objective – 'time'. This results in what are known as the 'iron triangle' of project management: quality, time and cost.

The relative importance of each objective will differ for different projects. Some aerospace projects, such as the development of a new aircraft, which impact on passenger safety, will place a very high emphasis on quality objectives. With other projects, for example a research project that is being funded by a fixed government grant, cost might predominate. Other projects emphasise time: for example, the organisation of an open-air music festival has to happen on a particular date if the project is to meet its objectives. In each of these projects, although one objective might be particularly important, the other objectives can never be totally forgotten.

> OPERATIONS PRINCIPLE
> *Different projects will place different levels of emphasis on cost, time, and quality objectives.*

Good objectives are those which are clear, measurable and, preferably, quantifiable. Clarifying objectives involves breaking down project objectives into three categories : the purpose, the end results and the success criteria. For example, a project that is expressed in general terms as 'improve the budgeting process' could be broken down into:

- Purpose: to allow budgets to be agreed and confirmed prior to the annual financial meeting.
- End result: a report that identifies the causes of budget delay, and which recommends new budgeting processes and systems.
- Success criteria: the report should be completed by 30 June, meet all departments' needs and enable integrated and dependable delivery of agreed budget statements. Cost of the recommendations should not exceed $200,000.

Project scope

The scope of a project identifies its work content and its products or outcomes. It is a boundary-setting exercise that attempts to define the dividing line between what each part of the project will do and what it won't do. Defining scope is particularly important when part of a project is being outsourced. A supplier's scope of supply will identify the legal boundaries within which the work must be done. Sometimes the scope of the project is articulated in a formal 'project specification'. This is the written, pictorial and graphical information used to define the output, and the accompanying terms and conditions. The project scope will also outline limits or exclusions to the project. This is critical, because perceptions of project success or failure often originate from the extent to which deliverables, limits and exclusions have been clearly stated and understood by all parties during the scoping phase.

Project strategy

The third part of a project's definition is the project strategy, which defines, in a general rather than a specific way, how the project is going to meets its objectives. It does this in two ways: by defining the phases of the project, and by setting milestones, and/or 'stagegates'. Milestones are important events during the project's life and often trigger payments to contractors. Stagegates are the decision points that allow the project to move onto its next phase. A stagegate often launches further activities and therefore commits the project to additional costs, and so on. At this stage of project planning, the actual dates for each milestone or stagegate are not necessarily determined. It is useful, however, to at least identify the significant milestones and stagegates, either to define the boundary between phases, or to help in discussions with the project's customer.

DIAGNOSTIC QUESTION

Is project management adequate?

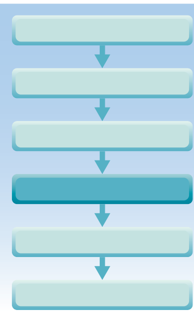

In order to coordinate the efforts of many people in different parts of the organisation (and often outside it as well), all projects need a project manager. Many of a project manager's activities are concerned with managing human resources. The people working in the project team need a clear understanding of their roles in the (usually temporary) organisation. Controlling an uncertain project environment requires the rapid exchange of relevant information with the project stakeholders, both within and outside the organisation. People, equipment and other resources must be identified and allocated to the various tasks. Undertaking these tasks successfully makes the management of a project a particularly challenging operations activity.

Project management skills

The project manager is the person responsible for delivering a project. She or he leads and manages the project team, with the responsibility, if not always the authority, to run the project on a day-to-day basis. They are special people. They must possess seemingly opposing skills. They must be able to influence without necessarily having authority, pay attention to details without

OPERATIONS PRINCIPLE
The activity of project management requires interpersonal as well as technical skills.

losing sight of the big picture, establish an open, communicative environment while remaining wedded to project objectives and have an ability to hope for the best but plan for the worst. It is a formidable role. Ideally, it involves leading, communicating, organising, negotiating, managing conflict, motivating, supporting, team building, planning, directing, problem solving, coaching and delegating.

In more formal terms, typical project manager responsibilities include the following:

- devising and applying an appropriate project management framework for the project
- managing the creation of the required deliverables
- planning, monitoring and controlling the project
- delegating project roles within agreed reporting structures
- managing project risks, including the development of contingency plans
- liaising with programme management (if the project is part of a programme)
- liaising with senior management to assure the overall direction and integrity of the project
- initiating corrective action where necessary through the reallocation of resources
- managing changes to project objectives or details
- reporting through agreed reporting lines on project progress and stage assessments
- identifying and obtaining support and advice required for the management of the project
- managing ongoing project administration
- conducting post-project evaluation to assess how well the project was executed
- preparing any follow-on action recommendations as required.

This wide variety of responsibilities requires a range of skills. Five characteristics, in particular, are seen as important in an effective project manager:[6]

1. Background and experience which are consistent with the needs of the project
2. Leadership and strategic expertise, in order to maintain an understanding of the overall project and its environment, while at the same time working on the details of the project
3. Technical expertise in the area of the project in order to make sound technical decisions
4. Interpersonal competence and the people skills to take on such roles as project champion, motivator, communicator, facilitator and politician
5. Proven managerial ability, in terms of a track record of getting things done.

EXAMPLE

Project failure is nothing new[7]

Don't think that large government projects that go wrong are a recent phenomenon. Project specification changes, bad communication, schedule delays and simple bad luck are nothing new and have always been able to bring down even the most high-profile projects. In 1628, the Vasa, the most magnificent warship ever built for the Royal Swedish Navy was launched, in front of an excited crowd. It had sailed less than a few thousand metres during its maiden voyage in the waters of the Stockholm harbour, when, suddenly, after a gun salute was shot in celebration, the Vasa heeled over and, as water gushed in through the gun ports, vanished beneath the surface killing 53 of the 150 passengers. Shocked officials were left questioning how such a disaster could happen.

Yet, as a project, the story of the Vasa displayed many of the signs of potential failure. When her construction began in 1625, the Vasa was designed as a small traditional warship, similar to many others previously built by the experienced shipbuilder Henrik Hybertsson. Soon after, the Swedish King Gustav II Adolphus, at that time fighting the Polish Navy in the Baltic Sea, started ordering a series of changes to the shape and the size of the warship, making its design much longer and bigger than originally envisaged. Also his spies informed the King that the Danes had started building warships with two gun decks, instead of the customary one. This would give the Danes a great advantage in terms of superior firepower from a longer distance.

From the battlefront the King mailed his order to add a second gun deck to the Vasa. The message caused consternation when it reached the shipbuilder several months later, but they attempted to comply with the change even though it caused wasteful reworking and complex patching up, as no one had ever seen or built such a revolutionary design before. Yet more pressure was put on the project when a catastrophic storm in the Baltic Sea, destroyed ten of the King's ships, making the commissioning of the Vasa even more urgent. Then, as a final piece of bad luck (especially for him), the shipbuilder, Hybertsson, died. Nevertheless, just before the ship's completion, a Navy representative Admiral Fleming conducted a stability test to assess the seaworthiness of the ship. Notwithstanding the strong signals of instability during the test, the Vasa was launched on its maiden voyage – with disastrous results for the King, for the Swedish Navy and for the passengers.

Managing matrix tensions

In all but the simplest project, project managers usually need to reconcile the interests of both the project itself and the departments contributing resources to the project. When calling on a variety of resources from various departments, projects are operating in a 'matrix management' environment, where projects cut across organisational boundaries and involve staff that are required to report to their own line manager as well as to the project manager. Figure 15.6 illustrates the type of reporting relationship that usually occurs in matrix management structures running multiple projects. A person in department 1, assigned part-time to projects A and B will be reporting to three different managers, all of whom will have some degree of authority over their activities. This is why matrix management requires a high degree of cooperation and communication between all individuals and departments. Although decision-making authority will formally rest with either the project, or departmental manager, most major decisions will need some degree of consensus. Arrangements need to be made that reconcile potential differences between project managers and departmental managers. To function effectively, matrix management structures should have the following characteristics.

- There should be effective channels of communication between all managers involved, with relevant departmental managers contributing to project planning and resourcing decisions.
- There should be formal procedures in place for resolving the management conflicts that do arise.
- Project staff should be encouraged to feel committed to their projects as well as to their own department.
- Project management should be seen as the central coordinating role, with sufficient time devoted to planning the project, securing the agreement of the line managers to deliver on time and within budget.

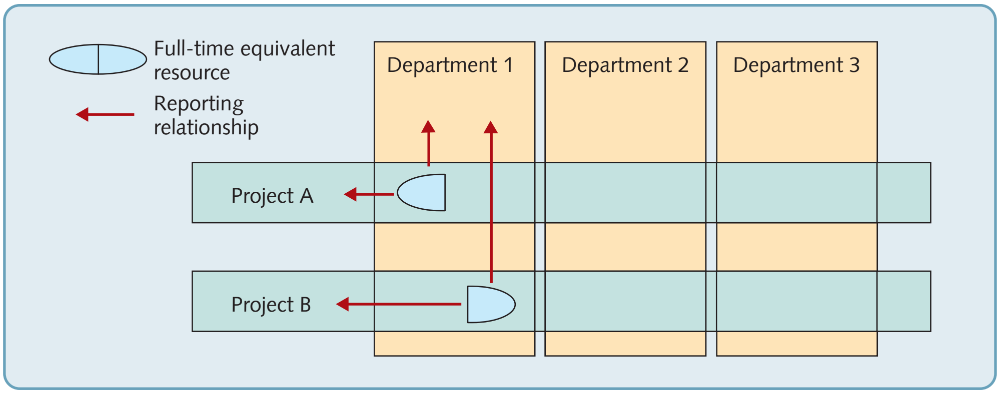

Figure 15.6 Matrix management structures often result in staff reporting to more than one project manager as well as their own department

DIAGNOSTIC QUESTION

Has the project been adequately planned?

All projects, even the smallest, need some degree of planning. The planning process fulfils four distinct purposes. It determines the cost and duration of the project, it determines the level of resources that will be needed, it helps to allocate work and to monitor progress and it helps to assess the impact of any changes to the project. It is a vital step at the start of the project, but it could be repeated several times during the project's life as circumstances change. This is not a sign of project failure or mismanagement. As discussed earlier, projects can and should be differentiated based on their characteristics – in our case, we examined three alternative approaches of volume–variety; uncertainty–complexity–scale; and novelty–technology–complexity–pace. And when managing particularly difficult projects, it is a normal occurrence to repeat planning throughout the project's life. In fact, later stage plans typically mean that more information is available, and that the project is becoming less uncertain. The process of project planning involves five steps shown in Figure 15.7.

OPERATIONS PRINCIPLE

A pre-requisite for project planning is some knowledge of times, resources and relationships between activities.

Identify activities – the work breakdown structure

Some projects are too complex to be planned and controlled effectively unless they are first broken down into manageable portions. This is achieved by structuring the project into a 'family tree' that specifies the major tasks or sub-projects. These in turn are divided up into smaller tasks until a defined, manageable series of tasks, called a work package, is arrived at. Each work package can be allocated its own objectives in terms of time, cost and quality, and can be assigned specific responsibility for its delivery. Typically, work packages do not exceed 10 days (though in practice they often do – see project planning example later), should be independent from each other, should belong to one sub-deliverable, and should constantly be monitored. The output from this is called the work breakdown structure (WBS). The WBS brings clarity and definition to the project planning process. It shows 'how the jigsaw fits together'. It also provides a framework for building up information for reporting purposes.

For example, Figure 15.8 shows the work breakdown structure for a project to design a new information interface (a website screen) for a new sales knowledge management system that is being installed in an insurance company. The project is a cooperation between the company's IT systems department and its sales organisation. Three types of activity will be necessary to complete the project, training, installation and testing. Each of these categories is further broken down into specific activities as shown in Figure 15.8.

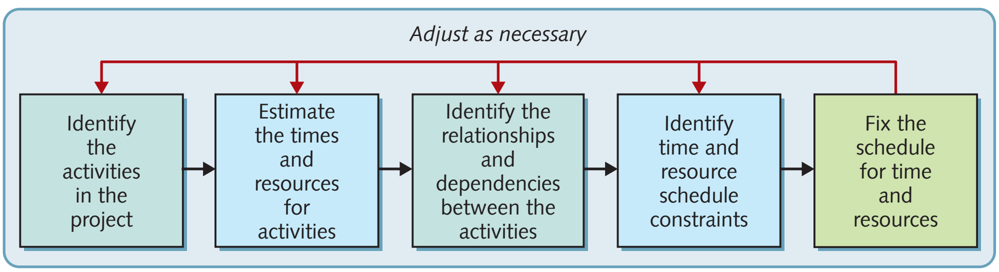

Figure 15.7 **Stages in the project planning process**

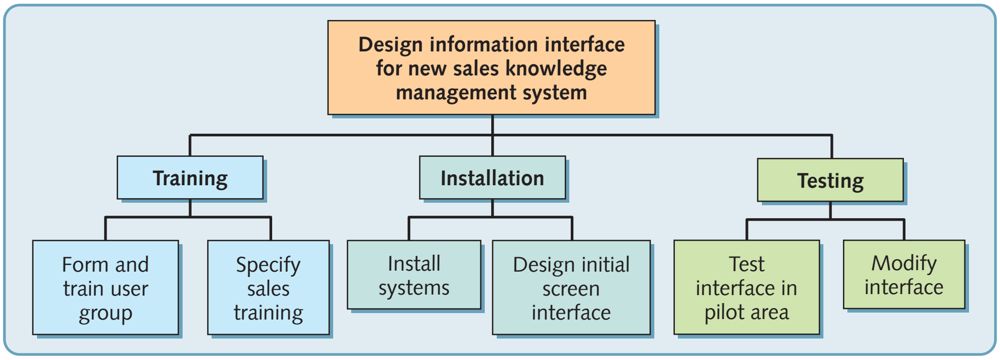

Figure 15.8 Work breakdown structure for a project to design an information interface for a new sales knowledge management system in an insurance company

Table 15.1 Time, resource and relationships for the sales system interface design project

Code	*Activity*	*Immediate predecessor(s)*	*Duration (Days)*	*Resources (Developers)*
a	Form and train user group	none	10	3
b	Install systems	none	17	5
c	Specify sales training	a	5	2
d	Design initial screen interface	a	5	3
e	Test interface in pilot area	b, d	25	2
f	Modify interface	c, e	15	3

Estimate times and resources

The next stage in planning is to identify the time and resource requirements of the work packages. Without some idea of how long each part of a project will take and how many resources it will need, it is impossible to define what should be happening at any time during the execution of the project. Estimates are just that – a systematic best guess, not a perfect forecast of reality. Estimates may never be perfect but they can be made with some idea of how accurate they might be. Table 15.1 includes time (in days) and resource (in terms of the number of IT developers needed) estimates for the sales system interface design project.

Probabilistic estimates

The amount of uncertainty in a project has a major bearing on the level of confidence that can be placed on an estimate. The impact of uncertainty on estimating times leads some project managers to use a probability curve to describe the estimate. In practice, this is usually a positively skewed distribution, as in Figure 15.9. More uncertainty increases the range of the distribution. The natural tendency of some people is to produce optimistic estimates, but these will have a relatively low probability of being correct because they represent the time that would be taken if everything went well. Most likely estimates have the highest probability of proving correct. Finally, pessimistic estimates assume that almost everything that could go wrong does go wrong. Because of the skewed nature of the distribution, the expected time for the activity will not be the same as the most likely time.

EXAMPLE

The BBC's Digital Media Initiative[8]

The BBC is one of the best-known broadcasters in the world, with an unrivalled reputation for the quality of some of its programmes. Sadly, its reputation for introducing new technology is less exemplary. Among its more spectacular failures was its Digital Media Initiative (DMI). It was an

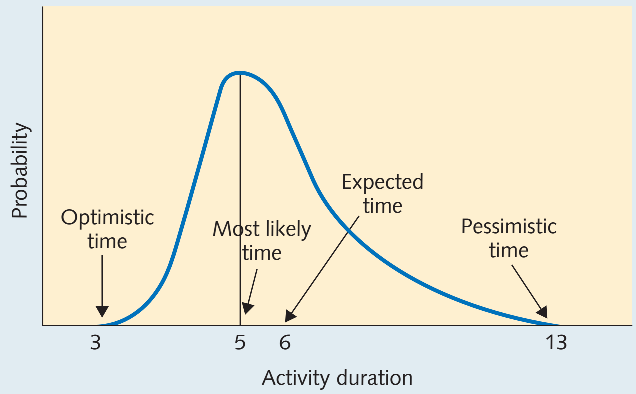

$$\text{Expected activity time} = t_o = \frac{t_o + 4t_1 + t_p}{6}$$

$$\text{Variance} = V = \frac{(t_p - t_o)^2}{36}$$

where

t_o = optimistic activity time

t_l = most likely activity time

t_p = pessimistic activity time

Figure 15.9 **Using probabilistic time estimates**

endeavour by the BBC to dispense with videotapes and create a kind of 'internal YouTube' of archive content that staff could access, upload, edit and then air from their computers. When the project was originally envisaged, creating a single TV programme could involve 70 individual video-handling processes. DMI was meant to halve that. The project cost almost £100 million and lasted five years before it was scrapped due to failures of governance and delayed delivery. The flaws in the technology were exposed during the BBC's coverage of the state funeral of Margaret Thatcher, a well-known ex-Prime Minister. The DMI was supposed to create a production system linked to the BBC's huge broadcasting archive, but instead of streamlining access to old video footage, video editors were unable to access archive footage to use in news reports from their computers in central London. They had to transport videotapes there using taxis and the underground network from the archive storage facility in North West London. Admitting that to continue with the project would be '*throwing good money after bad*' the BBC suspended its chief technology officer. One BBC manager called the DMI project '*the axis of awful*', while another said, '*The scale of the project was just too big, and it got out of hand*.' Tony Hall, the BBC's director general, said that off-the-shelf tools '*that simply didn't exist five years ago*', had now become available and they could do the same job as some elements of the DMI. Experts, commenting on the BBC DMI case, said, '*it is not the biggest or the worst IT project failure in the public or private sectors and, without organisations*' implementing measures to guard against them, it will almost certainly not be the last'. Others put many of the problems down to '*scope creep*' where users kept changing their requirements.

Identify the relationships and dependencies between the activities

The third stage of planning is to understand the interactions between different project work packages. All the work packages (or activities) that are identified as comprising a project will have some relationship with one another that will depend on the logic of the project. Some activities will, by necessity, need to be executed in a particular order. For example, in the

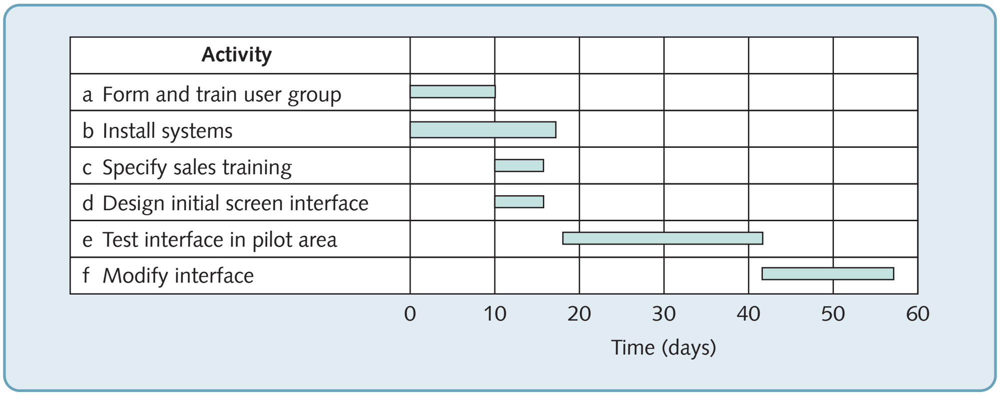

Figure 15.10 Gantt chart for the project to design an information interface for a new sales knowledge management system in an insurance company

construction of a house, the foundations must be prepared before the walls are built, which in turn must be completed before the roof is put in place. These activities have a dependent or series relationship. Other activities do not have any such dependence on each other. The rear garden of the house could probably be prepared totally independently of the garage being built. These two activities have an independent or parallel relationship.

In the case of the sales system interface design, Table 15.1 provided the basic information that enables the relationships between activities in the project to be established. It did this by identifying the immediate predecessor (or predecessors) for each activity. So, for example, activities a and b can be started without any of the other activities being completed. Activity **c** cannot begin until activity **a** has been completed, nor can activity **d.** Activity **e** can only start when both activities **b** and **d** have been completed, and activity **f** can only start when activities **c** and **e** have been completed.

Planning tools

Project planning is greatly aided by the use of techniques that help to handle time, resource and relationships complexity. The simplest of these techniques is the *Gantt chart* (or bar chart) that we introduced in Chapter 10. Figure 15.10 shows a Gantt chart for the activities that form the sales system interface project. The bars indicate the start, duration and finish time for each activity. The length of the bar for each activity on a Gantt chart is directly proportional to the calendar time, and so indicates the relative duration of each activity. Gantt charts are the simplest way to exhibit an overall project plan, because they have excellent visual impact and are easy to understand. They are also useful for communicating project plans and status to senior managers as well as for day-to-day project control.

As project complexity increases, it becomes more necessary to identify clearly the relationships between activities, and show the logical sequence in which activities must take place. This is most commonly done by using the *critical path method* (CPM) to clarify the relationships between activities diagrammatically. Though there are alternative methods of carrying out critical path analysis, by far the most common, and also the one used in most project management software packages, is the 'activity-on-node' (AoN) method. Figure 15.11 shows this for the sales team interface design project.

In the AoN representation, activities are drawn as boxes, and arrows are used to define the relationships between them. In the centre of each box is the description of the activity (in this case 'activity **a**', 'activity **b**', and so on). Above the description is the duration (D) of the activity (or work package), the earliest start time (EST) and earliest finish time (EFT). Below the description is the latest start time (LST), the latest finish time (LFT) and the 'float' (F) (the number of

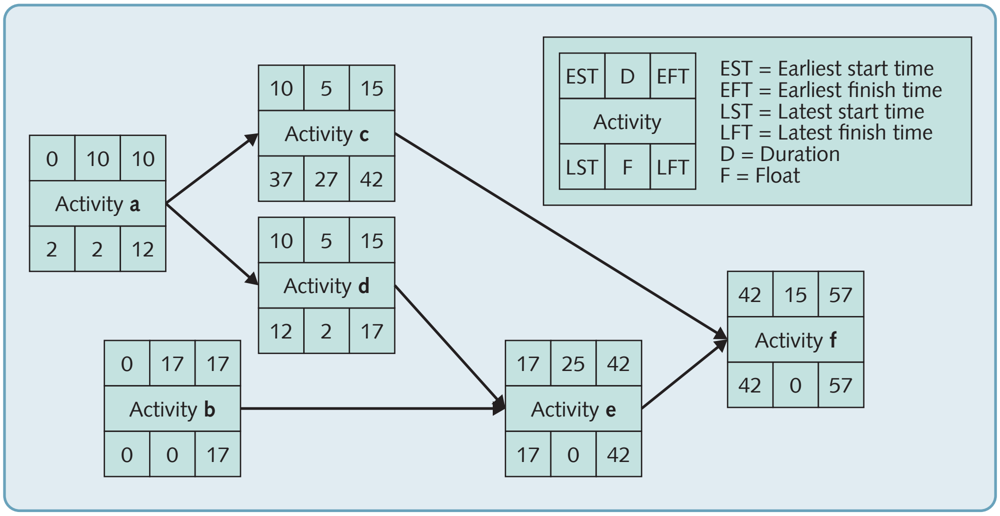

Figure 15.11 Critical path analysis (activity-on-node method) for the project to design an information interface for a new sales knowledge management system in an insurance company

extra days that the activity could take without slowing down the overall project. The diagram shows that there are a number of chains of events that must be completed before the project can be considered as finished (event 5). In this case, activity chains **a – c – f,** and **a – d – e – f,** and **b – e – f,** must all be completed before the project can be considered as finished. The longest (in duration) of these chains of activities is called the 'critical path' because it represents the shortest time in which the project can be finished, and therefore dictates the project timing. In this case **b – e – f** is the longest path and the earliest the project can finish is after 57 days.

Activities that lie on the critical path will have the same earliest and latest start times and earliest and latest finish times. That is why these activities are critical. Non-critical activities, however, have some flexibility as to when to start and finish. This flexibility is quantified into a figure that is known either as 'float' or 'slack'. So, activity **c,** for example, is only of 5 days duration and it can start any time after day 10 (when activity **a** is completed) and must finish any time before day 42 (when activities **a, b, c** and **d** are completed). Its 'float' is therefore $(42 - 10) - 5 = 27$ days (i.e. latest finish time minus earliest start time minus activity duration). Obviously, activities on the critical path have no float; any change or delay in these activities would immediately affect the whole project. In addition to a critical path (or network) diagram, the idea of float or slack can be shown diagrammatically on a Gantt chart, as in Figure 15.12. Here, the Gantt chart for the project has been revisited, but this time the time available to perform each activity (the duration between the earliest start time and the latest finish time for the activity) has been shown.

Identify time and resource schedule constraints

Once estimates have been made of the time and effort involved in each activity, and their dependencies identified, it is possible to compare project requirements with the available resources. The finite nature of critical resources – such as staff with special skills – means that they should be taken into account in the planning process. This often has the effect of highlighting the need for more detailed replanning.

The logic that governs project relationships, as shown in the critical path analysis (or network diagram), is primarily derived from the technical details, but the availability of resources may also impose its own constraints, which can materially affect the relationships between activities. Return to the sales system interface design project. Figure 15.13 shows the resource profile under two

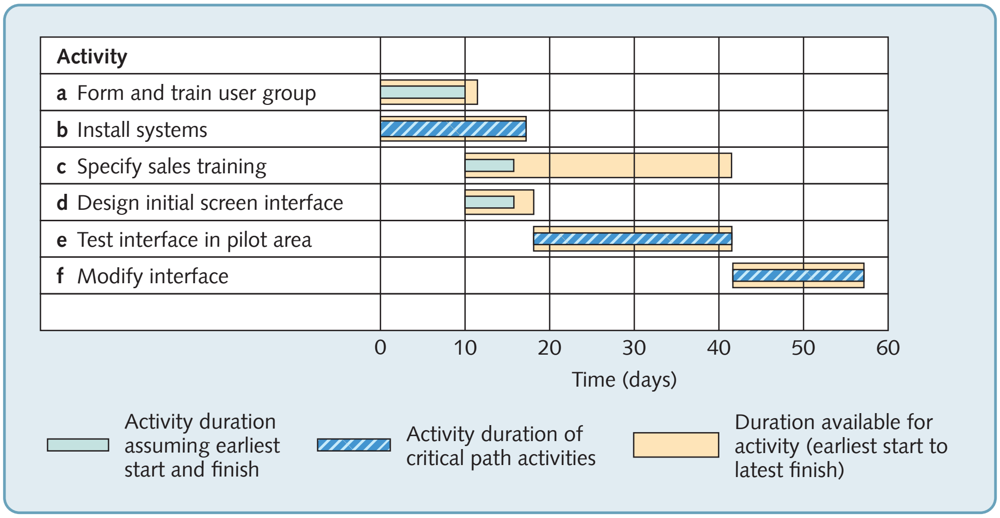

Figure 15.12 **Gantt chart for the project to design an information interface for a new sales knowledge management system in an insurance company with latest and earliest start and finish times indicated**

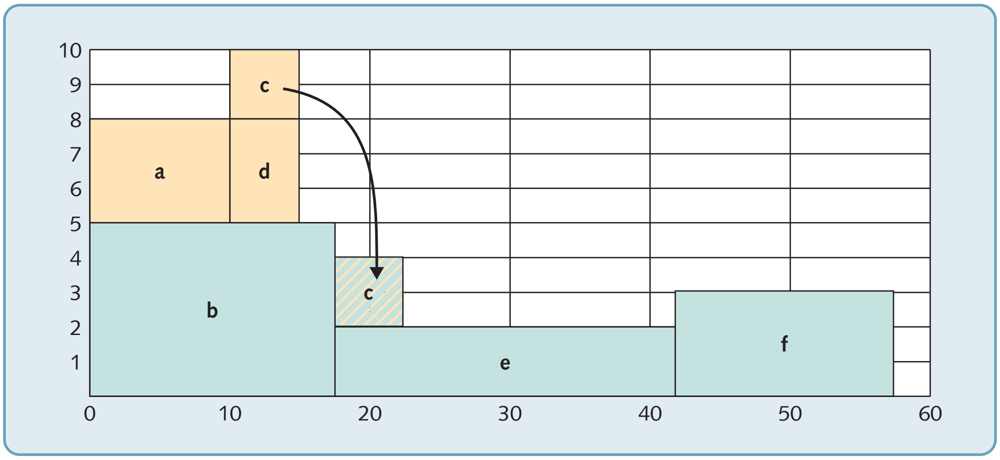

Figure 15.13 **Resource profiles for the sales knowledge system interface design, assuming that all activities are started as soon as possible, and assuming that the float in activity c is used to smooth the resource profile**

different assumptions. The critical path activities (**b – e – f**) form the initial basis of the project's resource profile. These activities have no float and can only take place as shown. However, activities **a, c** and **d** are not on the critical path, so project managers have some flexibility as to when these activities occur, and therefore when the resources associated with these activities will be required. From Figure 15.13, if one schedules all activities to start as soon as possible, the resource profile peaks between days 10 and 15 when 10 IT development staff are required. However, if the project managers exploit the float that activity c possesses and delays its start until after activity b has been completed (day 17), the number of IT developers required by the project does not exceed 8. In this way, float can be used to smooth out resource requirements or make the project fit resource constraints. However, it does impose further resource constrained logic on the relationship between the activities. So, for example, in this project moving activity c as shown in Figure 15.13 results in a further constraint of not starting activity c until activity b has been completed.

Fix the schedule for time and resources

Project planners should ideally have a number of alternatives to choose from. The one which best fits project objectives can then be chosen or developed. While it can be challenging to examine several alternative schedules, especially in very large or very uncertain projects computer-based software packages such as Bitrix24, Trello, 2-Plan PMS, Asana®, MS project and Producteev make critical path optimisation more feasible. The rather tedious computation necessary in network planning can relatively easily be performed by project planning models. All they need are the basic relationships between activities, together with timing and resource requirements for each activity. Earliest and latest event times, float and other characteristics of a network can be presented, often in the form of a Gantt chart. More significantly, the speed of computation allows for frequent updates to project plans. Similarly, if updated information is both accurate and frequent, such computer-based systems can also provide effective project control data.

Program evaluation and review technique (PERT)

While it is beyond the scope of this book to enter into much more detail of the various ways that critical path analysis can be made more sophisticated, programme evaluation and review technique (PERT) is worth noting as it represents an enrichment of the basic network approach. PERT, as it is universally known, originated in the planning and controlling of major defence projects in the US Navy, with its most spectacular gains in the highly uncertain environment of space and defence projects. The technique recognises that activity durations and costs in project management are not deterministic (fixed), and that probability theory can be applied to estimates, as was shown in Figure 15.9.

OPERATIONS PRINCIPLE

Probabilistic activity time estimates facilitate the assessment of a project being completed on time.

In this type of network, each activity duration is estimated on an optimistic, a most likely and a pessimistic basis, and the mean and variance of the distribution that describes each activity can be estimated as was shown in Figure 15.10 (see Table 15.2).

In this case the sum of the expected times for each of the activities on the critical path (**b – e – f**) is 58.17 days and the sum of the variances of these three activities is 6.07 days. From this, one can calculate the probability of the project overrunning by different amounts of time.

Table 15.2 PERT analysis for the sales system interface design project

Code	*Activity*	*Optimistic estimate*	*Most likely estimate*	*Pessimistic estimate*	*Expected time*	*Variance*
a	Form and train user group	8	10	14	10.33	1
b	Install systems	10	17	25	17.17	0.69
c	Specify sales training	4	5	6	5	0.11
d	Design initial screen interface	5	5	5	5	0
e	Test interface in pilot area	22	25	27	24.83	0.69
f	Modify interface	12	15	25	16.17	4.69

DIAGNOSTIC QUESTION

Is the project adequately controlled?

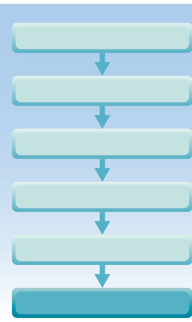

All the stages in project management described so far have taken place before the actual project takes place. Project control deals with activities during the execution of the project. Project control is the essential link between planning and doing.

The process of project control involves three sets of decisions about how to:

- monitor the project in order to check on its progress
- assess the performance of the project by comparing monitored observations of the project with the project plan
- intervene in the project in order to make the changes that will bring it back to plan.

Project monitoring

Project managers have first to decide what they should be looking for as the project progresses. Usually a variety of measures are monitored. To some extent, the measures used will depend on the nature of the project. However, common measures include current expenditure to date, supplier price changes, amount of overtime authorised, technical changes to project, inspection failures, number and length of delays, activities not started on time, missed milestone, and so on. Some of these monitored measures affect mainly cost, some mainly time. However, when something affects the quality of the project, there are also time and cost implications. This is because quality problems in project planning and control usually have to be solved in a limited amount of time.

Assessing project performance

A typical planned cost profile of a project through its life is shown in Figure 15.14. At the beginning of a project some activities can be started, but most activities will be dependent on finishing. Eventually, only a few activities will remain to be completed. This pattern of a slow start followed by a faster pace, with an eventual tail-off of activity holds true for almost all projects, which is why the rate of total expenditure follows an S-shaped pattern as shown in Figure 15.14, even when the cost curves for the individual activities are linear. It is against this curve that actual costs can be compared in order to check whether the project's costs are being incurred to plan. Figure 15.14 shows the planned and actual cost figures compared in this way. It shows that the project is incurring costs, on a cumulative basis, ahead of what was planned.

Intervening to change the project

If the project is obviously out of control in the sense that its costs, quality levels or times are significantly different from those planned, then some kind of intervention is almost certainly required. The exact nature of the intervention will depend on the technical characteristics of the project, but it is likely to need the advice of all the people who would be affected. Given the interconnected nature of projects – a change to one part of the project will have knock-on effects elsewhere – this means that interventions often require wide consultation. Sometimes

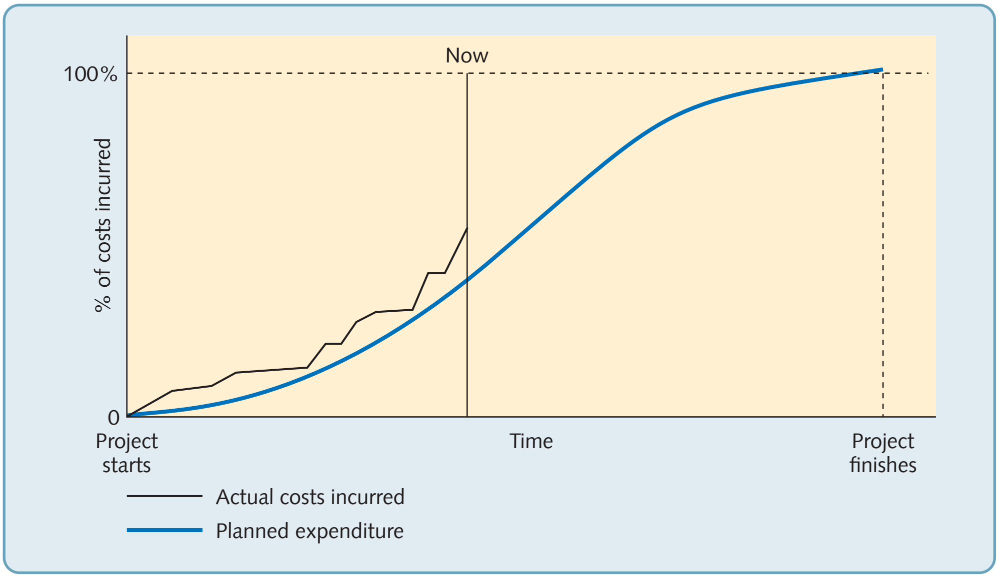

Figure 15.14 Comparing planned and actual expenditure

intervention is needed even if the project looks to be proceeding according to plan. For example, the schedule and cost for a project may seem to be 'to plan', but when the project managers project activities and cost into the future, they see that problems are very likely to arise. In this case, it is the trend of performance that is being used to trigger intervention.

Crashing, or accelerating, activities

Crashing activities is the process of reducing time spans on critical path activities so that the project is completed in less time. Usually, crashing activities incurs extra cost. This can be as a result of:

- overtime working
- additional resources, such as manpower
- sub-contracting.

Figure 15.15 shows an example of crashing a simple network. For each activity, the duration and normal cost are specified, together with the (reduced) duration and (increased) cost of crashing them. Not all activities are capable of being crashed; here activity e cannot be crashed. The critical path is the sequence of activities **a, b, c, e.** If the total project time is to be reduced, one of the activities on the critical path must be crashed. In order to decide which activity to crash, the 'cost slope' of each is calculated. This is the cost per time period of reducing durations. The most cost-effective way of shortening the whole project then is to crash the activity on the critical path, which has the lowest cost slope. This is activity a, the crashing of which will cost an extra €2,000 and will shorten the project by one week. After this, activity c can be crashed, saving a further two weeks and costing an extra €5,000. At this point all the activities have become critical and further time savings can only be achieved by crashing two activities in parallel.

> **OPERATIONS PRINCIPLE**
>
> ***Only accelerating activities on the critical path(s) will accelerate the whole project.***

The shape of the time–cost curve in Figure 15.15 is entirely typical. Initial savings come relatively inexpensively if the activities with the lowest cost slope are chosen. Later in the crashing sequence the more expensive activities need to be crashed and eventually two or more paths become jointly critical. Inevitably by that point, savings in time can only come from crashing two or more activities on parallel paths.

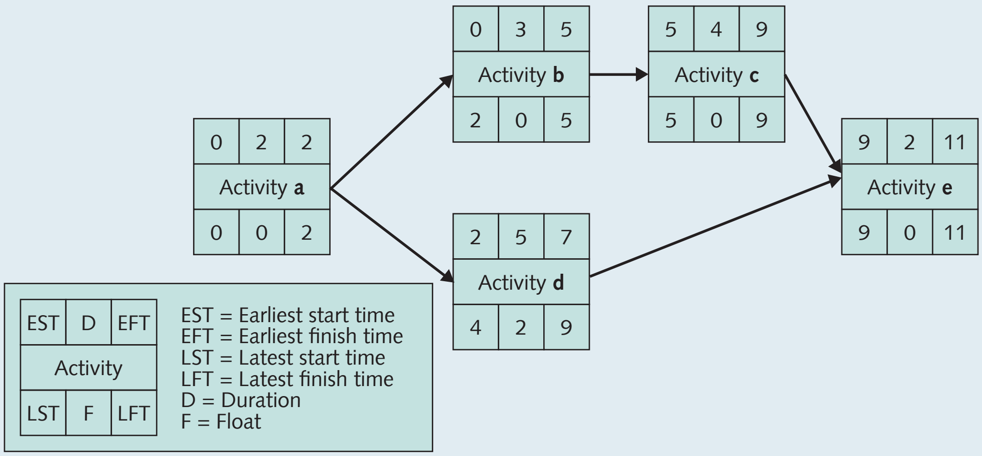

Activity	*Normal*		*Crash*		*Cost slope (€ 000/week)*
	Cost (€ 000)	*Time (weeks)*	*Cost (€ 000)*	*Time (weeks)*	
a	6	2	8	1	2
b	5	3	8	2	3
c	10	4	15	2	2.5
d	5	5	9	4	4
e	7	2	Not possible		–

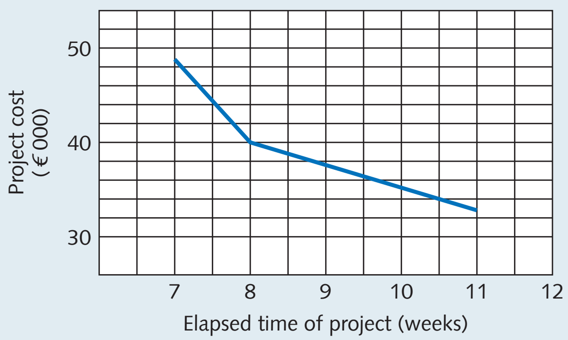

Figure 15.15 **Crashing activities to shorten project time becomes progressively more expensive**

Critical commentary

When project managers talk of 'time estimates', they are really talking about guessing. By definition, planning a project happens in advance of the project itself. Therefore, no one really knows how long each activity will take. Of course, some kind of guess is needed for planning purposes. However, some project managers believe that too much faith is put in time estimates. The really important question, they claim, is not how long will something take, but how long could something take without delaying the whole project. Also, if a single most likely time estimate is difficult to estimate, then using three, as one does for probabilistic estimates, is merely over analysing what are highly dubious data in the first place.

- The idea that all project activities can be identified as entities with a clear beginning and a clear end point and that these entities can be described in terms of their relationship with each other is another obvious simplification. Some activities are more or less continuous and evolve over time. For example, take a simple project such as digging a trench and laying a communications cable in it. The activity 'dig trench' does not have to be completed before the activity 'lay cable' is started. Only two or three metres of the trench need to be dug before cable laying can commence. This is a simple relationship, but one that is difficult to illustrate using critical path analysis. Also, if the trench is being dug in difficult terrain, the time taken to complete the activity, or even the activity itself, may change, to include rock-drilling activities for example. However, if the trench cannot be dug because of rock formations, it may be possible to dig more of the trench elsewhere – a contingency not allowed for in the original plan. So, even for this simple project, the original critical path or network diagram may reflect neither what will happen nor could happen.

SUMMARY CHECKLIST

- ☐ Do managers understand what projects are and the role of project management?
- ☐ Is the fundamental nature of the project understood in terms of volume–variety; scale–complexity–uncertainty; and novelty–technology–complexity–pace dimensions?
- ☐ Do managers understand the importance of identifying and managing stakeholders?
- ☐ Do managers have methods of stakeholder mapping and determining the nature of these stakeholders, in terms of their positive versus negative energy towards the project?
- ☐ Have all stakeholders been prioritised in terms of their relative power and interest?
- ☐ Has the project been well defined?
- ☐ Have the objectives of the project been defined, particularly in terms of the relative importance of cost, time and quality?
- ☐ Has the scope of the project been defined, including technical specifications, limits and exclusions?
- ☐ Has the overall strategy of the project been defined in terms of its overall approach, its significant milestones, and any decision gateways that may occur in the project?
- ☐ Have overall project management skills within the business been generally assessed?
- ☐ For this particular project, does the project manager have skills appropriate for the project's intrinsic degree of difficulty?
- ☐ Is sufficient effort being put into the project planning process?
- ☐ Have all activities been identified and expressed in the form of a work breakdown structure?
- ☐ Have all activity times and resources been estimated using the best possible information within the organisation?
- ☐ Is there sufficient confidence in the time and resource estimates to make planning meaningful?
- ☐ Have the relationships and dependencies between activities been identified and summarised in the form of a simple network diagram?
- ☐ Have project planning tools, such as critical path analysis or PERT, been considered for the project?
- ☐ Have potential resource and time schedule constraints been built into the project plan?
- ☐ Are there mechanisms in place to monitor the progress of the project?
- ☐ Have mechanism for intervening in the project to bring it back to plan been put in place?

CASE STUDY

United Photonics Malaysia Sdn Bhd

Introduction

Anuar Kamaruddin, COO of United Photonics Malaysia (UPM), was conscious that the project in front of him was one of the most important he had handled for many years. The number and variety of the development projects underway within the company had risen sharply in the last few years, and although they had all seemed important at the time, this one – the 'Laz-skan' project – clearly justified the description given it by the President of United Photonics Corporation, the US parent of UPM, *'. . . the make or break opportunity to ensure the division's long term position in the global instrumentation industry.'*

The United Photonics Group

United Photonics Corporation had been founded over 70 years ago (as the Detroit Gauge Company), a general instrument and gauge manufacturer for the engineering industry. By expanding its range into optical instruments, it eventually moved also into the manufacture of high-precision and speciality lenses, mainly for the photographic industry. Its reputation as a specialist lens manufacturer led to such a growth in sales that soon the optical side of the company accounted for about 60 per cent of total business and it ranked one of the top two or three optics companies of its type in the world. Although its reputation for skilled lens making had not diminished since then, the instrument side of the company came to dominate sales as the market for microchip-making equipment expanded.

UPM product range

UPM's product range on the optical side included lenses for inspection systems that were used mainly in the manufacture of microchips. These lenses were sold both to the inspection system manufacturers and to the chip manufacturers themselves. They were very high-precision lenses; however, most of the company's optical products were specialist photographic and cinema lenses. In addition, about 15 per cent of the company's optical work was concerned with the development and manufacture of 'one or two off' extremely high-precision lenses for defence contracts, specialist scientific instrumentation and other optical companies. The group's instrument product range consisted largely of electromechanical assembles with an increasing emphasis on software-based recording, display and diagnostic abilities. This move towards more software-based products had led the instrument side of the business towards accepting some customised orders. The growth of this part of the instrumentation had resulted in a special development unit being set up; the Customer Services Unit (CSU) who modified, customised, or adapted products for those customers who required an unusual product. Often CSU's work involved incorporating the company's products into larger systems for a customer.

Some years earlier United Photonics Corporation had set up its first non-North American facility just outside Kuala Lumpur in Malaysia. United Photonics Malaysia Sdn Bhd (UPM) had started by manufacturing subassemblies for Photonics instrumentation products, but soon had developed in a laboratory for the modification of United Photonics products for customers throughout the Asian region. This part of the Malaysian business was headed by T.S. Lim, a Malaysian engineer who had taken his post-graduate qualifications at Stanford and three years ago moved back to his native KL to head up the Malaysian outpost of the CSU, reporting directly to Bob Brierly the vice-president of development, who ran the main CSU in Detroit. Over the last three years, T.S. Lim and his small team of engineers had gained quite a reputation for innovative development. Bob Brierly was delighted with their enthusiasm. *'Those guys really do know how to make things happen. They are giving us all a run for our money.'*

The Laz-skan project

The idea for Laz-skan had come out of a project which T.S. Lim's CSU had been involved with. CSU had successfully installed a high-precision Photonics lens into a character recognition system for a large clearing bank. The enhanced capability that the lens and software modifications had given had enabled the bank to scan documents even when they were not correctly aligned. This had led to CSU proposing the development of a 'vision metrology' device that could optically scan a product at some point in the manufacturing process, and check the accuracy of up to 20 individual dimensions. The geometry of the product to be scanned, the dimensions to be gauged and the tolerances to be allowed, could all be programmed into the control-logic of the device. The T.S. Lim team were convinced that the idea could have considerable potential. The proposal, which the CSU team had called the Laz-skan project, was put forward to Bob Brierly, who both saw the potential value of the idea and was again impressed by the CSU team's enthusiasm. *'To be frank, it was their evident enthusiasm that influenced me as much as anything. Remember that the Malaysian CSU had only been in existence for two years at this time – they were a group of keen but relatively young engineers. Yet their proposal was well thought out and, on reflection, seemed to have considerable potential.'*

In the November following their proposal to Brierly, Lim and his team were allocated funds (outside the normal budget cycle) to investigate the feasibility of the Laz-skan idea. Lim was given one further engineer and a technician, and a three-month deadline to report to the board. In this time he was expected to overcome any fundamental technical problems, assess the feasibility of successfully developing the concept into a working prototype, and plan the development task that would lead to the prototype stage.

The Lim investigation

T.S. Lim, even at the start of his investigation, had some firm views as to the appropriate 'architecture' for the Laz-skan project. By 'architecture' he meant the major element of the system, their functions and how they related to each other. The Laz-skan system architecture would consider five major sub-systems: the lens and lens mounting, the vision support system, the display system, the control-logic software and the documentation.

T.S. Lim's first task, once the system's overall architecture was set, was to decide whether the various components in the major sub-systems would be developed in-house, developed by outside specialist companies from UPM's specifications, or bought in as standard units and if necessary modified in-house. Lim and his colleagues made these decisions themselves, while recognising that a more consultative process might have been preferable. *'I am fully aware that ideally we should have made more use of the expertise within the company to decide how units were to be developed. But within the time available we just did not have the time to explain the product concept, explain the choices, and wait for already busy people to come up with a recommendation. Also there was the security aspect to think of. I'm sure our employees are to be trusted but the more people who know about the project, the more chance there is for leaks. Anyway, we did not see our decisions as final. For example, if we decided that a component was to be bought in and modified for the prototype building stage, it does not mean that we can't change our minds and develop a better component in-house at a later stage.'* By February, T.S.'s small team had satisfied themselves that the system could be built to achieve their original technical performance targets. Their final task before reporting to Brierly would be to devise a feasible development plan.

Planning the Laz-skan development

As a planning aid the team drew up a network diagram for all the major activities within the project from its start through to completion, when the project would be handed over to manufacturing operations. A simplified activity-on-node (AoN) network is shown in Figure 15.16. The duration of all the activities in the project were estimated either by T.S. Lim or (more often) by him consulting a more experienced engineer back in Detroit. While he was reasonably confident in the estimates, he was keen to stress that they were just that – estimates.

1) The lens

The lens was particularly critical since the shape was complex and precision was vital if the system was to perform up to its intended design specification. T.S. Lim was relying heavily upon the skill of the group's expert optics group in Pittsburg to produce the lens to the required high tolerance. Since what in effect was a trial and error approach was involved in their manufacture, the exact time to manufacture would be uncertain. T.S. Lim realised this.

'The lens is going to be a real problem. We just don't know how easy it will be to make the particular geometry and precision we need. The optics people won't commit themselves even though they are regarded as some of the best optics technicians in the world. If the development goes wrong it could overrun substantially. It is a relief that lens development is not amongst the "critical path" activities.'

2) Vision support system

The vision support system included many components that were commercially available, but considerable engineering effort would be required to modify them. Although the development design and resting of the vision support system was complicated, there was no great uncertainty in the individual activities, or therefore the schedule of completion. If more funds were allocated to their development, some tasks might even be completed ahead of time.

3) The control software

The control software represented the most complex task, and the most difficult to plan and estimate. In fact, the software development unit had little experience of this type of work

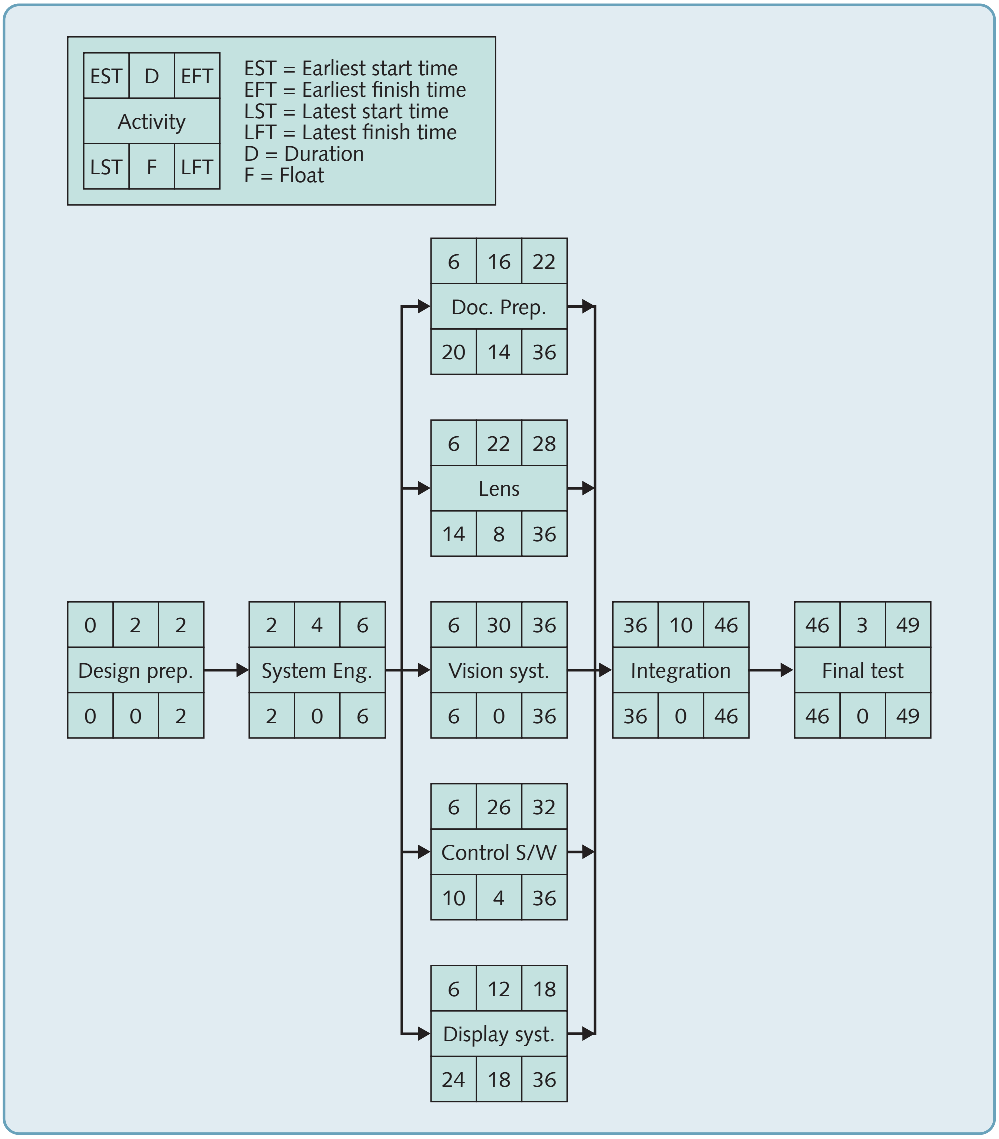

Figure 15.16 **Network diagram for the Laz-skan development**

but (partly in anticipation of this type of development) had recently recruited a young software engineer with some experience of the type of work that would be needed for Laz-skan. He was confident that any technical problems could be solved even though the system needs were novel, but completion times would be difficult to predict with confidence.

4) Documentation

A relatively simple sub-system, 'documentation' included specifying and writing the technical manuals, maintenance routines, on-line diagnostics and 'help desk' information. It was a relatively predictable activity, part of which was subcontracted to technical writers and translation companies in Kuala Lumpur.

5) Display system

The simplest of the sub-systems to plan, the display system, would need to be manufactured entirely out of the company and tested and calibrated on receipt.

Market prospects

In parallel with T.S. Lim's technical investigation, sales and marketing had been asked to estimate the market potential of Laz-skan. In a very short time, the Laz-skan project had aroused considerable enthusiasm within the function, to the extent that Halim Ramli, the Asian Marketing vice-president, had taken personal charge of the market study. The major conclusions from this investigation were:

1 The global market for Laz-skan type systems was unlikely to be less than 50 systems per year in 2 years' time, climbing to more than 200 per year after 5 years.
2 The volume of the market in financial terms was more difficult to predict, but each system sold was likely to represent around US$300,000 of turnover.
3 Some customisation of the system would be needed for most customers. This would mean greater emphasis on commissioning and post-installation service than was necessary for UPM's existing products.
4 Timing the launch of Laz-skan would be important. Two 'windows of opportunity' were critical. The first and most important was the major world trade show in Geneva in almost exactly a year's time (the following April). This show, held every two years, was the most prominent showcase for new products such as Laz-skan. The second related to the development cycles of the original equipment manufacturers that would be the major customers for Laz-skan. Critical decisions were generally taken in the fall. If Laz-skan was to be incorporated into these companies products it would have to be available within 18 months (the September of the following year).

The Laz-skan go ahead

At the end of February UPM considered both the Lim and the Ramli reports. In addition, estimates of Laz-skan's manufacturing costs had been sought from George Hudson, the head of instrument development. His estimates indicated that Laz-skan's operating contribution would be far higher than the company's existing products. The board approved the immediate commencement of the Laz-skan development through to prototype stage, with an initial development budget of US$4.5 million. The objective of the project was to, '*. . . build three prototype Laz-skan systems to be 'up and running' for April 2006*'.

The decision to go ahead was unanimous. Exactly how the project was to be managed provoked far more discussion. The Laz-skan project posed several problems. First, engineers had little experience of working on such a major project. Second, the crucial deadline for the first batch of prototypes meant that some activities might have to be been accelerated, an expensive process that would need careful judgment. Finally, no one could agree either whether there should be a single project leader, which function he or she should come from, or how senior the project leader should be. Anuar Kamaruddin knew that these decisions could affect the success of the project, and possibly the company, for years to come.

QUESTIONS

1 Who do you think should manage the Laz-skan development project?
2 What are the major dangers and difficulties that will be faced by the development team as they manage the projects towards its completion?
3 What can they do about these dangers and difficulties?

APPLYING THE PRINCIPLES

Some of these exercises can be answered by reading the chapter. Others will require some general knowledge of business activity and some might require an element of investigation. **Model answers for the first two questions are available on the student companion website.** For model answers to the other questions in this section, please ask your tutor/lecturer.

1 Revisit the two examples towards the beginning of the chapter describing the type of projects undertaken by Disney Imagineering and the Crossrail project. Using the three methods of distinguishing between projects that are described in the chapter (their volume and variety characteristics; their scale, complexity and degree of uncertainty; their novelty, complexity, nature of technology and 'pace'), analyse the two types of project.

2 The activities, their durations and precedences for designing, writing and installing a bespoke computer database are shown in Table 15.3 below. Draw a network diagram (activity-on-node) for the project and calculate the fastest time in which the operation might be completed.

Table 15.3 Bespoke computer database activities

Activity	*Duration (weeks)*	*Activities that must be completed before it can start*
1 Contract negotiation	1	–
2 Discussions with main users	2	1
3 Review of current documentation	5	1
4 Review of current systems	6	2
5 Systems analysis (a)	4	3,4
6 Systems analysis (B)	7	5
7 Programming	12	5
8 Testing (prelim)	2	7
9 Existing system review report	1	3,4
10 System proposal report	2	5,9
11 Documentation preparation	19	5,8
12 Implementation	7	7,11
13 System test	3	12
14 Debugging	4	12
15 Manual preparation	5	11

3 '*Funding comes from a variety of sources; to restore the literally irreplaceable buildings we work on. We try to reconcile historical integrity with commercial viability, and rely on the support of volunteers. So we need to involve all stakeholders all the way through the project*.' (Janine Walker, chief project manager, Happy Heritage, a not-for-profit restoration organisation.) Her latest project was the restoration of a 200-year-old 'poorhouse' as a visitor attraction, originally built to house local poor. Janine's team drew up a list of stakeholders and set out to win them over with their enthusiasm for the project. They invited local people to attend meetings, explained the vision and took them to look round the site. Also before work started Janine took all the building staff on the same tour of the site as they had taken other groups and the VIPs who provided the funding. '*Involving the builders in the project sparked a real interest in the project and the archaeological history of the site. Often they would come across something interesting, tell the foreman who would involve an archaeologist and so preserve an artefact that might otherwise have been destroyed. They took a real interest in their work, they felt involved*.' (a) Who do you think would be the main stakeholders for this project? (b) How might not involving them damage the project, and how would involving them benefit the project?

4 Some (even relatively experienced) project managers neglect stakeholders in the project management process, preferring to 'manage them at a distance', rather than allow them to interfere with the project. Others argue that the benefits of stakeholder management are too great to ignore. It has been suggested that emphasising the *responsibilities,* as well as the rights, of project stakeholders can moderate many of the risks of stakeholder involvement. If you were in charge of drawing up a general list of the rights and responsibilities of project stakeholders in an information technology company, what would you include?

5 In the oil industry, project teams are increasingly using virtual reality and visualisation models of offshore structures that allow them to check out not only the original design, but any modifications that have to be made during construction.

(a) Why do you think a realistic picture of a completed project helps the process of project management?

(b) Why are such visualisations becoming more important?

6 Examine this simple domestic project. The project definition is to make 'breakfast in bed' consisting of a boiled egg, toast and orange juice, using the minimum staff resources and time, and to a high quality (egg freshly boiled, warm toast, etc.). The project is to start in the kitchen at 6.00 a.m., and finish in the bedroom. The activities involved in the project, resources and times, are shown in Table 15.4. Draw a chart that shows your recommended start and finish times for each activity.

Table 15.4 Time and resources estimates for a 'breakfast-in-bed' project

Activity	*Effort (person-min)*	*Duration*
a) Butter toast	1	1
b) Pour orange juice	1	1
c) Boil egg	0	4
d) Slice bread	1	1
e) Fill pan with water	1	1
f) Bring water to boil	0	3
g) Toast bread	0	2
h) Take loaded tray to bedroom	1	1
i) Fetch tray, plates, cutlery	1	1

Notes on chapter

1 Reiss, G. (1996) *Programme Management Demystified,* E. & F.N. Spon.

2 Sources include: Walt Boyes, W. (2011) 'David Van Wyk shares the Disney Imagineering project management process', controlglobal.com, 29 June; Mark T. Hoske, M.T. (2011) 'What do Walt Disney Imagineering and NASA space travel have in common? Engineering inspiration', *Control Engineering,* 22 September, http://www.controleng.com/

3 Sources include: *Bechtel report,*(2017) 'Better train lines for London commuters', http://www.bechtel.com/projects/crossrail-london/, accessed 1 May; Marrs, C. (2016) 'Crossrail: on time and on budget, is this how to get a major infrastructure project right?', *Civil Service World,* 8 February.

4 Sources include: raconteur.net,(2015) *Project management,*2 August; World Health Organization malaria programme http://www.who.int/topics/malaria/en/

5 Shenhar, A.J. and Dvir, D. (2007). *Reinventing Project Management: The Diamond Approach to Successful Growth and Innovation,* Harvard Business School Press.

6 Weiss, J.W. and Wysocki, R.K. (1992) *Five-Phase Project Management: A Practical Planning and Implementation Guide,* Addison-Wesley.

7 This example was written and kindly supplied by Mattia Bianchi, Department of Management and Organization, Stockholm School of Economics

8 Gilick, B. (2014) 'The BBC DMI project – what went wrong?', *Computer Weekly,* 5 February.

TAKING IT FURTHER

There are hundreds of books on project management. They range from the introductory to the very detailed, and from the managerial to the highly mathematical. Here are three general (as opposed to mathematical) books which are worth a look.

Cole, R. and Scotcher, E. (2015) ***Brilliant Agile Project Management: A Practical Guide to Using Agile, Scrum and Kanban, Pearson Business.*** *A practical and modern take on project management.*

Lock, D. (2013) ***Project Management, Routledge.*** *A classic text.*

Maylor, H. (2018) ***Project Management (5th edn), Pearson.*** *A good basic text on the subject.*

Index

'This highly influential textbook is and always will be on the operations management billboard.'
Professor Pär Åhlström, Stockholm School of Economics, Sweden

'This book provides valuable insights into how to design, plan and improve operations. The diagnostic questions in each chapter are excellent for guiding the reader to the essence of each topic.'
Dr Erik Soepenberg, Senior Lecturer, Hanze University of Applied Sciences, Groningen, the Netherlands

'An insightful and engaging blend of theory and practice that demonstrates the fundamental importance of operations and process management to today's organisations.'
Nigel Spinks, Henley Business School, University of Reading, UK

Written by the best-selling authors in their field, the fifth edition of *Operations and Process Management* inspires a critical and authoritative mastery of the core principles and processes that are fundamental to successfully managing business operations.

Approaching the subject from a managerial perspective, this text provides clear, comprehensive and concise coverage of the principles and practice of operations and process management.

Features

- ***New!*** Over 20 new and updated examples and cases demonstrate how operations work in practice.
- ***New!*** Restructured treatment of the scope and structure of operations activities emphasise the importance of supply network issues.
- ***New!*** Fresh coverage of such topical issues as behavioural operations, valuable operations capabilities, surge pricing, 'digital twins', the 'circular economy' and operations risk.
- ***New!*** Extended coverage of concepts such as sustainability, creativity and innovation in the design process, value stream mapping, SIPOC analysis and the 'line of sight' in internal process networks.
- Cases and examples drawn from a wide range of industries and regions provide an insight into the reality of managing operations and give practical understanding of the material.
- The use of 'diagnostic logic chains' allows effective evaluation and improvement of operations and processes.

This is *the* ideal text for students taking any practically-oriented MBA, Executive MBA, or executive course in Operations and Process Management.

www.pearson-books.com

Pearson

ISBN 978-1-292-17613-0

9 781292 176130 >