



# **Building on Knowledge**

**Developing Expertise,  
Creativity and Intellectual  
Capital in the Construction  
Professions**

**David Bartholomew**



**WILEY-BLACKWELL**

A John Wiley & Sons, Ltd., Publication

This edition first published 2008  
© 2008 Blackwell Publishing Ltd

Blackwell Publishing was acquired by John Wiley & Sons in February 2007.  
Blackwell's publishing programme has been merged with Wiley's global Scientific,  
Technical, and Medical business to form Wiley-Blackwell.

*Registered office*

John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester, West Sussex,  
PO19 8SQ, United Kingdom

*Editorial offices*

9600 Garsington Road, Oxford, OX4 2DQ, United Kingdom  
2121 State Avenue, Ames, Iowa 50014-8300, USA

For details of our global editorial offices, for customer services and for information  
about how to apply for permission to reuse the copyright material in this book  
please see our website at [www.wiley.com/wiley-blackwell](http://www.wiley.com/wiley-blackwell).

The right of the author to be identified as the author of this work has been asserted  
in accordance with the Copyright, Designs and Patents Act 1988.

All rights reserved. No part of this publication may be reproduced, stored in a  
retrieval system, or transmitted, in any form or by any means, electronic, mechanical,  
photocopying, recording or otherwise, except as permitted by the UK Copyright,  
Designs and Patents Act 1988, without the prior permission of the publisher.

Wiley also publishes its books in a variety of electronic formats. Some content that  
appears in print may not be available in electronic books.

Designations used by companies to distinguish their products are often claimed as  
trademarks. All brand names and product names used in this book are trade names,  
service marks, trademarks or registered trademarks of their respective owners. The  
publisher is not associated with any product or vendor mentioned in this book. This  
publication is designed to provide accurate and authoritative information in regard to  
the subject matter covered. It is sold on the understanding that the publisher is not  
engaged in rendering professional services. If professional advice or other expert  
assistance is required, the services of a competent professional should be sought.

ISBN: 978-1-4051-4709-5

Library of Congress Cataloging-in-Publication Data is available

A catalogue record for this book is available from the British Library.

Set in 10 on 12.5 pt Avenir by SNP Best-set Typesetter Ltd., Hong Kong  
Printed in Singapore by Fabulous Printers Pte Ltd



# Contents

<i>Preface</i>	vii
<i>Acknowledgements</i>	xi
<b>Part One Foundations</b>	<b>1</b>
<b>1 Introduction</b>	<b>3</b>
Paradoxical professionals	5
New context, new issues	9
What is in this book	12
<b>2 Knowledge at Work</b>	<b>15</b>
How we learn	15
What makes an expert	19
Varieties of knowledge	22
Putting the pieces together	27
<b>3 Strategic Frameworks</b>	<b>34</b>
Starting points	34
Frameworks for thinking	35
Finding conviction	41
<b>4 The Challenges of Change</b>	<b>44</b>
Why initiatives fail	44
Difficulty is normal	59
<b>5 Leadership and Other Roles</b>	<b>61</b>
Action starts where the buck stops	61
Practical leadership	63
Other roles	70
Knowledge-conscious management	78
<b>6 Knowledge Audit and Beyond</b>	<b>79</b>
Finding square one	79
Audit techniques	83

From audit to action plan	89
Putting plans into practice	92
<b>Part Two Tools and Techniques</b>	<b>95</b>
<b>7 The Knowledge-Friendly Office</b>	<b>97</b>
Environments matter	97
Designing the knowledge-friendly office	99
Workplaces for teams	102
<b>8 Expanding Networks</b>	<b>106</b>
It's not what you know . . .	106
Help from IT	108
Designing networking tools	111
<b>9 Learning from Peers</b>	<b>119</b>
See one, do one, teach one	119
Mentoring in different contexts	122
<b>10 Learning from Practice</b>	<b>128</b>
Practice: the invisible lab and unsung teacher	129
Windows of opportunity	130
Foresight: learning from invention	131
Hindsight: learning from mistakes – and success	135
Choosing cases	144
<b>11 Communities of Practice</b>	<b>146</b>
Encouraging enthusiasts	146
Creating communities	148
<b>12 Organisational Memory</b>	<b>151</b>
The indispensability of the written word	151
Deciding what to record, and how	154
Capturing knowledge	156
Documenting knowledge	159
Software frameworks	165
<b>13 Personal Knowledge Management</b>	<b>176</b>
Equipment for the mind gym	176
Developing personal expertise	177
Building a bionic memory	179
<b>14 Synergies</b>	<b>181</b>
IT-enabled synergies: networking directories, knowledge bases and business systems	181
Creating and sharing knowledge: foresight, hindsight and knowledge bases	184

Multiple synergies: communities of practice, knowledge bases and mentoring	185
<b>Part Three Knowledge Management in Practice</b>	<b>187</b>
<b>15 Introduction to the Case Studies</b>	<b>189</b>
The case studies	189
Recurring patterns	190
<b>16 Aedas</b>	<b>193</b>
Starting points	194
MIS	194
Aedas Studio	197
Knowledge audit	199
Emerging knowledge systems	202
Commentary	204
<b>17 Arup</b>	<b>206</b>
Starting points	206
Projects	207
Future	210
Commentary	211
<b>18 Broadway Malyan</b>	<b>213</b>
Starting points	213
Business Process	215
Who's Who	218
Contact database	219
Induction process	221
Commentary	222
<b>19 Buro Happold</b>	<b>223</b>
Starting points	224
The prototype	224
The final design	225
Assessing the results	228
Commentary	229
<b>20 Edward Cullinan Architects</b>	<b>231</b>
Starting points	233
Knowledge strategy	235
Commentary	241
<b>21 Feilden Clegg Bradley</b>	<b>244</b>
Starting points	245
Hindsight reviews	246

Yellow Pages	248
Knowledge base	249
Commentary	251
<b>22 Penoyre &amp; Prasad</b>	<b>254</b>
Starting points	255
The R&D database	256
The knowledge bank	257
Lessons learned	260
Commentary	261
<b>23 Whitbybird</b>	<b>263</b>
1: Identifying knowledge systems and assets	264
2: Selecting a subset for audit	265
3: Choosing audit methods	265
4: Designing the questionnaire	266
5: Testing and refining the questionnaire	266
6: Conducting the survey	266
7: Analysing the results	267
Commentary	268
<b>24 WSP</b>	<b>269</b>
Starting points	269
Technical coordinator workshops	271
Commentary	272
<b>25 Case Studies on Foresight and Hindsight</b>	<b>273</b>
Amicus Group	275
BAA	277
BP/Bovis Global Alliance	279
Buro Happold	282
Lattice Property	284
<b>Epilogue</b>	<b>289</b>
<b>26 Where Next for Knowledge Management?</b>	<b>291</b>
Web 2.0	292
The Semantic Web	293
Developments in psychology and the science of human relations	294
Insights from neuroscience	294
<i>Further Reading</i>	296
<i>Index</i>	301



## Preface

My interest in knowledge began over 25 years ago when I was responsible for directing the UK's national programme of research on solar energy. As results started to roll in it became clear that, though some of the ideas belonged to the future, others deserved to be taken up immediately: they could make buildings cheaper to run and nicer to live and work in, without costing anything. Only the designs would need to change. We published the research reports, but we soon found that hardly anybody in the construction industry reads research reports. How could we get these new ideas across to them? How, indeed, did knowledge in general flow from research into practice?

When I looked into past innovations, I found that it could take up to 20 years for new ideas to spread throughout a whole industry (and construction was not uniquely slow). I did not want to wait that long. Even after new ideas reached one part of a company they often took a long time to become common knowledge. How could I speed up the process? I became fascinated by knowledge and how it flows around, between organisations and inside them. I discovered a lot about how ideas emerge and practical know-how develops; that some kinds of knowledge can be communicated easily in writing, some only with difficulty, and some not at all; the importance of tacit knowledge (but I had not read Polanyi's book, so I did not know there was a name for it) and the special magic of face-to-face knowledge transfer; and many other things. And this mysterious stuff called knowledge has been part of my professional life ever since.

This book was inspired by a series of three research projects about knowledge in construction – principally in architectural practices, engineering consultancies and client organisations – which I initiated and led between 1998 and 2005. They focused respectively on the use of IT to make information more easily and quickly accessible, on learning from project experience, and on sharing knowledge within organisations – between them, all the main processes involved in what we have

come to call 'knowledge management'. We wanted to discover how organisations can improve quality, avoid endlessly reinventing wheels and repeating mistakes, reduce risk, become more creative, make working life more enjoyable, and improve their bottom lines.

Looking back, four things stand out: that understanding of knowledge management has developed enormously in the past decade, thanks largely to the accumulation of evidence from practice; that the same principles, tools and techniques have emerged as central to success in all kinds of organisation (including the principle that *implementation* needs to be tailored sensitively to the organisational context); that, despite these commonalities, professional services organisations have features that pose special problems for managing knowledge; and that, despite the many books – some of them excellent – that have been published on the subject, business leaders *still* find it hard to discover what to do. This is an attempt to fill some of the gap.

The practitioners I have worked with tell me they find the existing literature variously too academic, too didactic, too specialised, too abstract, too much concerned with the alien world of big corporations, and simply too extensive. Much of it also makes knowledge management sound alarmingly complicated. Software vendors, by contrast, claim that knowledge management is simply a matter of buying their (hugely expensive) 'solutions'. This is seductive, but it has long since been exposed as misleading at best and often a quick road to disappointment. Well-informed boards are now wisely sceptical.

This is, therefore, a consciously *non-academic*, *undidactic*, wide-ranging and relatively short book, which looks at knowledge management from a practical and specifically professional services perspective. There is, for example, no chapter reviewing the history of research or thought in knowledge management, and I have not tried to relate the various tools and techniques to academic theories. And although, like all books on knowledge management, this one has roots in a wide variety of ideas from management guru Peter Drucker (who coined the term 'knowledge worker' around 50 years ago) onwards, I have cited authorities for only a few of them. Many appear anyway to have arisen independently in several places. That is not surprising in a field where the best evidence for most propositions is simply personal experience. To adapt the inscription on Christopher Wren's tomb in St Paul's Cathedral: *lector, si argumentum requiris, circumspice*.<sup>1</sup>

It is not simply a collection of recipes, either. The trouble with recipes alone is that they give the cook no help when the meal needs to be adapted for different ingredients, equipment or tastes, because

---

<sup>1</sup> Reader, if you seek the justification, look around you.



they fail to explain *why* certain things work and others do not. Since the basic recipes of knowledge management always have to be adapted to suit organisations' individual circumstances, some understanding of underlying causalities is vital. This I have tried to provide. Nevertheless, well-tried recipes are invaluable as a starting point, so they are here too.

I have focused particularly on the needs of consultancy practices and repeat clients – architects, engineers, surveyors, and clients in sectors such as education, health, government, retail and the utilities – but much of the book should be equally relevant to the professional aspects of contracting. It does not, though, address transactional aspects of construction such as e-tendering and the handling of project documentation, which are essentially exercises in managing data and information, not knowledge.

It is probably too much to hope that people in other professional services will get past the title, but if they do look further they will find much that is relevant to them, too. Most professional services organisations share characteristics that differentiate them from the large manufacturing and commercial corporations which dominate the knowledge management literature. Most of them are small; their work is usually project based rather than process based; their staff are relatively homogeneous in educational background and function; management structures are flat; and, often, the managers are also the owners, and carry on their professional work alongside management. These all have important consequences for managing knowledge, which are discussed here.

For those who want to explore the subject further, I have included a list of selected books and articles as further reading.

Finally, I have kept explanations as simple as I can; knowledge management involves many interrelated issues and some quite subtle ideas, but it need not be difficult to understand. Professionals such as architects, consulting engineers and doctors have always done many of the things it involves, and the challenge is not to do something fundamentally new, but to do it more consciously, with more understanding, and so more effectively. It is not rocket science. If one has enough of the basic ideas clear in one's mind, the implications emerge naturally: deciding what to do just takes effort to work things through, and the determination and patience to act on the answers. And it is worth the effort: as Drucker said, 'The basic economic resource . . . is and will be knowledge.'

None of my research would have been possible without funding from the UK Department of Trade and Industry. I am deeply indebted to all the organisations that have been my active partners in it over the years, and to my consultancy clients; without the lessons I have learned working with them to address the challenges of real-world

knowledge management this book would not exist. I am especially grateful to Peter Oborn of Aedas, Chris Askew and Bill Gething of Feilden Clegg Bradley, Ashraf Michail of the BP/Bovis Global Alliance, Colin Rice of Edward Cullinan Architects, Andrew Cripps of Buro Happold, and Adrian Burton of Broadway Malyan. I cannot thank them and many others adequately for their openness about the sometimes untidy realities of practice, their willingness to take time out of pressured lives to debate ideas and try them out, their unfailing encouragement, and their kindness. Finally, I am eternally grateful to my wife Marion for casting a critical eye on drafts and for her understanding and tolerance. Nevertheless, despite all their invaluable contributions, any faults and shortcomings in this book are mine and mine alone.



# Acknowledgements

I would like to thank Aedas, Broadway Malyan, Buro Happold, Edward Cullinan Architects, Feilden Clegg Bradley, Penoyre & Prasad and Whitbybird for allowing me to use the screenshots, photographs and diagrams that appear in their case studies.

I am also grateful to Arup and Colin Rice for allowing me to include the personal accounts of knowledge management initiatives which appear in the Arup and Edward Cullinan Architects case studies respectively.





# **Part One**

## **Foundations**



# Chapter One

## Introduction

A *New Yorker* cartoon of a few years ago shows an elderly man being introduced to a group of young staff, all looking up from their laptops: 'For those of you who don't know Mr Ingham – he's our institutional memory.' It is a neat encapsulation of a predicament that many organisations face today: the loss of accumulated expertise and dilution of corporate ethos as baby boomers retire and greenhorns flood in. And it is a reminder of deeper truths, too: that knowledge lives in people's brains, not in computers, and much of it can only be shared face to face, if at all. As the leading management thinker of the 20th century, Peter Drucker, put it: 'Knowledge is between two ears, and only between two ears.'

The loss of organisational memory and capability when long-serving staff leave is just one consequence of the general difficulty of sharing knowledge, and particularly the practical know-how that accumulates with experience. When knowledge is locked up in individual brains and local teams, unshared – as most of it is, in most organisations – wheels are reinvented, old mistakes are repeated, misunderstandings create new ones, and good practice stubbornly fails to spread. In professional services, practices that fail to pool their knowledge find economies of scale elusive, and that growth brings less competitive advantage than it should. Often, it seems to do little more than create a federation of small practices that share overheads.<sup>1</sup>

The scale of the waste from reinvention alone can be surprising. When I polled the staff of a large and highly successful architectural practice in a recent knowledge audit, only 25% thought they spent less than 10% of their time reinventing wheels, and 37% thought they spent over 20%. The average guess was 18%: that is typical. People's estimates of the time they spend looking for information are usually similar. The total effect in wasted time, lower quality and lost profit is

---

<sup>1</sup> One of the reasons why boutiques continue to prosper alongside giants.

considerable, and there are other prices to pay in missed opportunities to increase quality, reduce risk, and improve in other ways.

Not long ago, managers just shrugged their shoulders at all this – if they thought about knowledge at all – but the effect on corporate performance is becoming harder to ignore in an increasingly demanding and competitive world. At the same time, the trend towards larger, more dispersed and more complex organisations, and higher labour mobility, is causing knowledge to fragment more than ever and making it harder to share it. The Mr Inghams might be able to pass on their knowledge to half a dozen people round a table, but not to hundreds or thousands, spread across a country or across the globe.

Growing awareness of the value of sharing knowledge is the main reason for the proliferation of corporate knowledge bases, skills directories, communities of practice and other tools and techniques designed to make it flow around more freely. But these are not enough. Brains are not just passive repositories of knowledge; they create it by absorbing new experience and reshaping and extending old knowledge to accommodate it. Every time we tackle a new challenge, whether it is a hard-fought game of chess, a tricky design problem, or a meeting with a difficult client, we have an opportunity to learn how to do better on the next occasion we meet a similar situation. That is what turns the theoretical knowledge we acquire at university into practical competence, and develops junior staff into respected seniors, the most able into experts, and run-of-the mill firms into industry leaders. But few people outside psychology faculties think consciously about the processes involved, and as a result most personal and organisational learning is subconscious, haphazard, and more or less inefficient. That is just as wasteful as poor knowledge-sharing. Research has shown convincingly that the top experts and the top firms are those who accumulate the most experience, and learn most effectively from it.

When firms compete in a free market it is the differences between them that make one succeed more than another; the common factors merely define the baseline for entry. That makes the knowledge that comes from experience particularly valuable: its uniqueness makes it a key differentiator, whereas other knowledge is available to everyone for the price of a journal subscription or a course fee. Toyota's competitors can hire engineers from the same universities, buy the same books, and even tour its factories, but they do not have access to its unique experience, and few have learned so much from their own, or shared what they have learned so effectively. It is the knowledge that Toyota has accumulated in its workforce's heads, in company documents, in patents and in other forms – what Thomas A. Stewart called 'intellectual capital' – that has made it the most profitable volume car maker in the world.



To prosper in the 21st century, organisations of all kinds will need to become much better both at creating new intellectual capital and at using what they already have. This will require two things: *understanding* of what knowledge is and how it flows around, and active *management* of the processes of learning, sharing and the accumulation of corporate knowledge – in other words, knowledge management.

## Paradoxical professionals

You would expect professional services organisations to be among the first to embrace knowledge management (or 'KM'); after all, knowledge is their stock in trade, and their staff and what they know are their largest asset. But no: with the notable exception of management consultancies (who were among the first), many of them have barely started. Who were the early adopters? The US Army, Toyota, Ford, Canon, Siemens, Chevron, BP . . . all organisations with huge assets of other kinds. To understand the paradox we need to look at how knowledge management has developed. That is worth doing because it shows why the time is ripe for professional services such as architecture, engineering, surveying and medicine to follow their lead, and why simply copying what they do will not work. It turns out that there are good reasons why the early adopters were fertile ground when professional services were not. Fortunately, it can be smart to be a late adopter: knowledge management is harder than it looks, and it helps to be able to learn from other people's successes and failures.

Knowledge management has had a meteoric rise. Before about 1995 the term was almost unknown, though some of the central ideas were already around under names such as 'organisational learning' and 'the learning organisation'. Today, it is familiar in the boardrooms of all kinds of organisation, across the world. The number of academic papers on the subject quadrupled between 1995 and 1997 and again by 1999, and books and articles on both theory and practice began to proliferate at the same time. Nonaka and Takeuchi's seminal *The Knowledge-Creating Company: How Japanese companies create the dynamics of innovation* reached the bookshops in 1995, and Stewart's *Intellectual Capital: The new wealth of organisations* appeared a couple of years later. Today a search for 'knowledge management' in Amazon.com books produces over 9000 hits; a general search on Google produces over 9 million.

Numerous organisations have taken up the idea and reported successes, often crediting it with major improvements in productivity and capability. As early as 1997 the Chief Executive of BP, John Browne, told a *Harvard Business Review* interviewer that improving learning

and knowledge-sharing had generated \$4 billion worth of permanent improvements in his company over the previous five years. When the Economist Intelligence Unit surveyed senior executives worldwide in 2005 and asked 'Which of the following areas of activity offer the greatest potential for productivity gains over the next 15 years?', knowledge management was the most popular choice by a wide margin. Assessing the changes likely in the global economy, industry and corporate structures over the same period, the EIU identified KM as one of the five principal trends, and concluded that improving the productivity of knowledge workers through technology, training and organisational change would be the major boardroom challenge of the next 15 years.

But there is another side to this rosy picture of progress, success and promise. Bain & Company have polled business executives almost every year since 1993 to see how widely various management tools are used, and how satisfied people are with them. By 2006 nearly 70% of the organisations surveyed reported using KM, with more planning to start in 2007, but only 17% reported being 'extremely satisfied' with it, and 16% were 'dissatisfied'. It ranked in the bottom 20% of tools for average satisfaction – as it has done every year since it was first included in Bain's survey. Satisfaction, of course, is a measure of the gap between expectation and achievement, and low satisfaction might only reflect unrealistically high expectations. That probably is a factor, but other evidence – and my own experience – suggests that low achievement *certainly* is. Booz Allen Hamilton estimated some years ago that only one KM programme in six achieves 'very significant' business impact in its first two years, half achieve 'small but important' benefits, and the remaining third are essentially failures. I suspect that little has changed since, despite the flood of advice in papers, books and conferences.

So is knowledge management a runaway success and a strategic priority for late adopters such as professional services, or is it a classic case of the emperor's new clothes – a deception nobody dares expose? I think it is something of both: a strategic priority and a success when realistic expectations and effective implementation coincide, but a disappointment when they do not. And it is too often made out to offer more than it really can, and to be easier to implement than it really is. To understand what it has to offer professional services we need to look beyond the generalisations of international, cross-industry surveys and consider what people mean when they talk about 'knowledge management', and why they continue to have such high hopes for it.

KM is a highly elastic concept, and it means very different things to different people. Software companies sell shrink-wrapped applications as 'knowledge management solutions' (none of them are!), and consultants and academics have described it in terms such as 'making

the best use of the knowledge the organisation has got', 'the capacity to take effective action', 'about how to get people to work smarter', and even 'not the management of knowledge'. In practice, 'knowledge management systems' often turn out to be little more than old information management systems rebranded with a fashionable name, or a collection of procedures and IT tools that hardly anyone uses. With such a wide variety of usages, making sweeping judgements about it is like making judgements about transport without distinguishing between cars, boats and planes, or between what is being carried, and where. Further, the fuzziness of the concept makes it difficult for managers to form a clear vision of it, what it entails, or what to expect of it, let alone implement it successfully. And it is hardly surprising that many initiatives fall short of high aspirations such as 'making the best use of the knowledge the organisation has got'. We shall consider later what knowledge management can usefully mean in professional practice.

Despite its ambiguity, it is not hard to see why the idea of KM took off when it did and in the industries where it did, and why people still have such high hopes for it despite its mixed success in practice. Several key factors coincided for the first time in the 1980s and 1990s, and together they made the importance of knowledge in business clearer than ever before, and provided both the inspiration and the tools to do something about it:

- *Intangible assets* such as knowledge, patents and brands became the largest components of corporate value. The Brookings Institute has estimated that, as recently as 1982, over 60% of the market capitalisation of companies in the S&P 500 index was based on tangible assets such as factories, machinery and stocks; by 1992 the proportion had fallen to under 40%, and by 2002 it was less than 15%. The balance – today over 90% of value – is based on intangibles. The rise of companies such as Microsoft made the trend obvious to everyone during the 1990s, and acute business leaders were not slow to recognise its implications.
- *Globalisation* put pricing pressure on manufacturers, and at the same time showed the West how much Japanese companies were benefiting from their close attention to knowledge. By the 1990s Japanese industry had become a force in a range of major industries, making well-designed products with a production efficiency and quality that Western competitors struggled (and mostly failed) to match. Manufacturers used to dismissing Japanese goods as derivative and cheap-and-cheerful found that customers increasingly saw brands like Sony, Canon and Honda as premium

options, worth premium prices. Not surprisingly, Nonaka and Takeuchi found many eager readers when *The Knowledge-Creating Company* showed how much of their success was based on a culture of continuous learning and widespread knowledge-sharing. Combining Japanese production methods with cheaper labour in Korea, and later elsewhere, turned the competitive screw even further, and made it imperative for Western companies to adopt similar techniques.

- *Quality* became an imperative, too. I do not know why customers lost patience with faulty products, but they did. Perhaps it was just that the Japanese had proved that high quality was possible, or maybe it was a reaction to changes in manufacturing methods that made repair disproportionately expensive. Governments took advantage of the new possibilities to tighten regulatory standards for food quality, hygiene, waste disposal, energy efficiency, health and safety, and various other aspects of operations and products. Where regulation was not feasible, they cajoled. To improve standards in the construction industry, for example, the UK government sponsored a report on *Rethinking Construction* that lambasted it for endemic cost escalation, time overruns and defects, and called for 'radical improvement' in quality and efficiency. Together, higher customer expectations and tighter regulation made faults and mistakes matter more than ever before.
- *Growing size and geographic dispersion* meant that informal, intuitive communication ceased to work in many companies. As the Chief Knowledge Officer of Ernst & Young is said to have remarked: 'In the old days we used to yell down the hall "Has anyone done this before?", but you can't yell down a hallway of 75 000 people.' – especially when they are spread across a continent, or even a city.
- *Management styles changed*. The trend away from command and control styles of management towards flatter structures required knowledge as well as authority to be shared more widely.
- *Publications* such as *The Knowledge-Creating Company* and HBR's interview with John Browne, 'Unleashing the power of learning', brought three crucial elements together for the first time in a style that business leaders could understand and apply: an intellectual foundation for thinking about corporate knowledge, persuasive evidence of the impact that learning and knowledge-sharing could have on business performance, and practical tools for making it happen.

- *Personal computers became universal for professionals, and the Internet established universal standards for data exchange. Together, these provided the technical means for people to communicate at a distance more freely than ever before, and for vast quantities of information to be stored and retrieved quickly and easily, from anywhere, by multiple users simultaneously.*

These origins go a long way towards explaining why large, mostly manufacturing, corporations were the first to adopt knowledge management: they felt the pressures of changing business conditions first and most strongly, and they could relate to the early success stories. At the same time, they explain why professional services have lagged behind. They were sheltered from the greatest pressures (it is difficult to outsource the design of a school or treatment for a broken leg to China, and harder to compare competing architects than cars), and other pressures, such as rising customer expectations and challenging regulation and performance targets, were weaker and generally later to arrive in their markets. In future the differences look like becoming much less.

A McKinsey survey in 2006 asked respondents what single factor contributed most to increasing competitive pressure on their industry. 'Improved capabilities of competitors' – in other words, better knowledge or better talent – came top, chosen by 25%, followed by 'more low-cost competitors' (23%). Ten per cent chose 'growing size of competitors', 8% 'regulatory changes' and 5% 'rising consumer awareness and activism'. These are all as recognisable in contexts such as construction and medicine as in other industries: faster learning and making better use of existing knowledge are rapidly becoming universal imperatives. The Economist Intelligence Unit was surely right to conclude that knowledge management will be one of the principal trends in affecting business through to 2020 – and nowhere more so than in professional services.

## New context, new issues

Even though they are increasingly subject to similar competitive pressures, professional services still differ from manufacturing companies in many ways, and will continue to do so. Expectations of KM and the way it is approached need to differ too. One of the key lessons from the past 10–15 years is that although the underlying principles of organisational learning and of knowledge-sharing apply everywhere, and many of the same basic tools and techniques can be used, the details of their implementation need to be tailored sensitively to the organisational context in order to succeed. Mies van der Rohe's

famous dictum that 'God is in the details' is just as apt for knowledge management as it is for his minimalist architecture. We shall consider the implications of this in later chapters, but it is worth pausing to review a few of the characteristic differences between professional services organisations and other industries, and their consequences. Unique rather than mass-replicated products, managers who also own the business and earn fees, project working and an ethos of individual autonomy all have implications for knowledge management, and on the whole they tend to make it more difficult to implement. These differences are a further reason for its late adoption in most professional services, and they need to be confronted to make it succeed.

Most industrial and commercial organisations develop products and then replicate them essentially identically and in large numbers – cars, TVs, PCs, socks, steel bars, barrels of oil, tonnes of aggregate, insurance policies, retail transactions, train journeys, you name it – whereas professional services organisations typically deal in one-offs such as buildings, medical treatments, and consultancy projects. This difference has several consequences.

The most significant is that volume replication multiplies the value of improvements, particularly in operational efficiency and product quality, and creates the possibility of big wins. Even one new idea, or the transfer of a good idea from one factory, office or shop to others, might repay the annual cost of a company's KM programme.

The scale of potential benefits can easily justify substantial investment in seeking improvements to individual products and processes. A structured programme of learning and knowledge-sharing at BP focused on oil refinery refurbishments, for example, cut direct costs by 20%, reduced the time they took by 9 days, and produced a longer-lasting result – a total saving of nearly \$10 million in each refurbishment, potentially repeated every 4–5 years and multiplied by around 20 refineries worldwide. Wins like this make both a strong business case for KM and good stories that can be a great help in convincing the indifferent and the sceptical. It is more difficult to justify generous investment in KM, and to motivate staff to make it work, in professional services, where the benefits are typically indirect, diffuse and largely unquantifiable, and big wins are almost impossible.

The role overlap between ownership, management and revenue-earning that is common in professional services is another factor that tends to make progress with knowledge management more difficult than it is in industries where they are separate. Its effects are particularly evident in medium-sized firms where ownership is shared relatively evenly between a dozen or more working partners or directors. Overlapping ownership and management puts decisions in the hands of people whose personal income is much more directly affected by short-term profitability than it is in quoted companies, and it may make investments in company-wide initiatives dependent on

consensus between a dozen or more people. That is bad news for activities like knowledge management, which offer benefits, however considerable, that are hard to pin down and may take years to realise, in return for immediate costs, however small.

Two of the central tenets of behavioural economics (which won Daniel Kahneman a Nobel prize in 2002) are that most people are loss-averse – they will forgo the possibility of substantial gains in order to avoid losses, and put more effort into avoiding a loss than into securing a gain – and that they put undue weight on near-term events and too little on far-off ones in making decisions. Even after an initial decision has been made to invest in KM, role overlaps can be a continuing obstacle to progress. The principal cost of knowledge management is in staff time, and even when intentions are good it can be hard for people at all levels – and particularly management – to wrench themselves away from more enjoyable, revenue-earning activities. This is an instance of a widespread management problem that Stanford professors Jeffrey Pfeffer and Robert Sutton christened the ‘knowing–doing gap’, and we shall return to it later. Further, when firms operate more like a collection of independent baronies than a unified organisation, as is not uncommon in professional services, a local equity-sharing director unconvinced by knowledge management can completely block progress on his patch. In an environment like this, even appointing a dedicated knowledge manager is unlikely to make much difference. In a discretionary, non-fee-earning activity and without either professional standing or equity his position is too weak.

Dealing in one-offs almost inevitably necessitates project working, another characteristic that distinguishes most professional services from other industries: design the building, complete the assignment, treat the patient, and move on to the next. The cessation of revenue from each project when it finishes, the variation between them, and the creative professional’s inner drive to try something new even when repetition might be more economic, all lead to a disinclination to look back systematically at completed projects in order to learn from them, let alone to make any effort to share lessons learned. Looking back costs money, a sacrifice of personal time, or both, and the lessons may be irrelevant in the next project. This is completely different from a typical manufacturing situation, where there is a conscious effort to make each new product an improvement on its predecessor, and to cut the cost of producing it, by identifying product weaknesses and process inefficiencies, finding ways to eliminate them, and mining competitors’ products for good ideas.

All these obstacles can be overcome by leaders and managers prepared to make difficult decisions: accept the possibility of a small short-term reduction in income; make any necessary financial investments; delegate in order to clear personal time for knowledge



management; give staff budgets for KM activities; make activities such as project reviews happen. But there are other obstacles that are less amenable to managerial determination. Professionals such as architects, consulting engineers and doctors are educated to expect considerable autonomy, and they are apt to believe that six or more years studying their discipline in university and in post-degree training has provided all the knowledge and skills they need. They are often reluctant to believe that anyone else can know better than they do, and strongly resistant to anything they see as interfering with their professional independence or creative freedom. Few professions have any tradition of looking elsewhere for ideas when people believe their existing knowledge is adequate.

The consequence is that many professionals search out information and advice only when they have to, and most tend to regard knowledge resources as a last rather than a first resort. Evidence-based medicine has only recently been accepted by doctors, and architects still show little inclination towards evidence-based design. Attitudes like these are far from unknown in other industries, but they are most deeply entrenched in the professions. A radical increase in learning and knowledge-sharing in an environment like this requires deep cultural change, and that poses a major challenge for business leaders who want their firms to use knowledge better.

Professional services organisations that have been late in adopting knowledge management, then, have not been perverse, but they would be perverse to delay much longer. As the management theorist Karl Sveiby has put it:

Managers often have an unconscious and tacit mindset that is coloured by the values and the common sense of the industrial age. To see another world, they need to try to use a conscious mindset such as the knowledge perspective.

There is an overwhelming case for making KM a strategic priority: in the short to medium term to improve competitiveness, and in the longer term as a prerequisite for survival. There is much that can be learned from the way in which other industries have taken it up over the past 10–15 years, but professional services differ from them in ways that make blindly copying their approaches, tools and techniques unlikely to succeed; they need to be adapted to suit the different environment. And knowledge initiatives will stand or fall largely on the clear thinking and determination of leaders and managers.

## What is in this book

This book has been written principally for partners, directors and managers (all 'managers' from now on, unless the distinctions are



important) in architectural, engineering, surveying and property consultancies who recognise the importance of organisational learning and knowledge-sharing for their future success, and want their firms to be better at both. Despite the focus on construction, I hope the issues it discusses will strike chords, and the ideas it presents will be helpful, for managers in other professional services as well, in both the public and private sectors.

It is intended equally for readers who have got no further than putting knowledge management on the 'to do' list, for those who are struggling to create a KM strategy or to make a knowledge initiative work, for those who want to overhaul existing tools and processes that no longer seem fit for purpose, and for those who want to improve further processes that already work well. Fundamentally, of course, these positions are all the same. Learning and knowledge-sharing are as old as the human race, and every organisation today has informal working practices, formal procedures and IT systems designed to assist them in one way or another. Only entirely new organisations have the luxury of starting with a clean sheet.

This book does *not* address the handling of operational documents such as correspondence, contracts, schedules, specs or drawings, or business information such as time sheets, personnel records and accounts. These contribute only indirectly to knowledge, and the specialised software that is designed to store them, make them readily accessible, enforce version control and so on (excellent as it may be for its purpose) has little relevance to the management of knowledge and the creation and use of intellectual capital.

This first part, *Foundations*, goes on in Chapter 2 to set the scene by reviewing knowledge, learning, knowledge management, and what they mean in a professional services context. Chapter 3 discusses how the aspirations and operational focus of an organisation define priorities for learning and knowledge-sharing. Chapter 4 addresses an issue that many books on knowledge management ignore, but which seems to me to be among the most crucial: why knowledge initiatives so often disappoint or fail entirely. Chapter 5 discusses the crucial importance of leadership in achieving success, and the other roles that need to be filled. Finally, Chapter 6 turns to practical details and explains how to use a knowledge audit to establish the status quo and set objectives for a knowledge initiative (whether aimed at radical change or minor improvement), and how to use the results to develop an action plan.

Part Two, *Tools and Techniques*, discusses the processes and IT tools that are most likely to be useful in professional services organisations. Chapters 7–14 deal respectively with workspace design, social networking software to help people with questions find people with answers, mentoring, processes for learning at the start and end of

projects ('foresight' and 'hindsight'), communities of practice (CoPs), the role of written knowledge and the software tools associated with it, personal knowledge management, and the relationships and synergies between them all. These chapters draw on experience accumulated over the past 15 years or so with the various tools and techniques in many kinds of organisation across the world, and discuss their strengths and weaknesses and how they can be tailored to suit the particular needs of professional services. Several chapters go into specific practical detail, but they are not recipes to be followed slavishly; rather, the detail is included to help readers visualise more clearly what the various tools and processes entail, and to provide a starting point for thinking creatively about them.

Part Three, *Knowledge Management in Practice*, describes some of the things that over a dozen of the most successful and managerially innovative companies in construction have done to improve their learning and knowledge-sharing. These are based on two research projects carried out between 2001 and 2005 in which I had the privilege of working closely with and advising them as they variously developed knowledge management strategies and implemented and tested new processes and tools. Most of the firms involved are professional practices, either architects or consulting engineers, but they also include the UK's largest airport operator (BAA), the BP/Bovis Global Alliance, a leading housing association, and others. They all started from different positions, and they followed a remarkable variety of paths. I am grateful to all the firms represented for their willingness to let me accompany them on their journeys, learn with them, and publish the details of what they did (and do) so that others can learn too from their difficulties and successes.

The *Epilogue* speculates on how organisational learning and knowledge-sharing might develop in the future.

# Chapter Two

## Knowledge at Work

We all think we know what knowledge is. It is such a pervasive part of life, and we say 'I know' so often without anyone asking for further explanation, that we rarely pause to consider what we really mean. But trying to manage an organisation's knowledge with only common usage as a guide is like trying to manage its finances with only a hazy idea of what money is – a recipe for disappointment, albeit with less immediately painful consequences. To be successful, knowledge management needs to be informed by a clearer understanding of the nature of knowledge, how it is created, and how people and organisations learn. Without that, managers are faced with a cascade of seemingly unanswerable questions when they try to choose tools and processes to match an organisation's particular needs, to get the details of their design and implementation right, or to take knowledge appropriately into account in other aspects of management. Why can't we just get all our experts to write down what they know? Isn't knowledge just information? What's special about face-to-face communication? What makes an expert? How can we make our new joiners productive more quickly? Why is it so difficult to transfer good practice from our London office to Newcastle? Why should we invite so many people to project reviews – the project leader can do it, can't he? Why should co-locating the design team help reduce project overruns? Why isn't it a waste of time for people to chat round the coffee machine?

### How we learn

Research has shown that human infants develop a remarkably sophisticated understanding of causality, mechanics and other people's minds years before they acquire the language to talk about them. They quickly learn, for example, that for one brick to push another it has to be in contact with it, and they start to use objects as tools to

extend their reach. They are surprised when an experimenter makes something happen that appears to violate causality. These are almost uniquely human abilities; even chimps fail to learn that a stick can be used to pull things towards them. We acquire fundamental knowledge like this by interacting with the world, observing it and thinking about what we have seen – by learning from experience – and that continues to be one of our most valuable sources of knowledge throughout life. It is, of course, the basis of all science and technology, too.

Causality shapes the way we understand the world, and the way we structure much of our knowledge. We start early to ask ‘why’ questions (as parents know to their cost), and the urge to find causal explanations continues to be strong throughout life. When observation fails to provide them we often invent them, and we may even invent entirely false memories to support them – all entirely unconsciously, of course. Memory is anything but the mental video recording we tend to imagine it to be: research has shown that we re-create our past every time we recall it, often slightly differently. And our behaviour is not controlled as much by the conscious, rational part of our brain as we usually assume. Our unconscious determines much more of our behaviour than most of us like to believe – it does well over 90% of our thinking, according to recent research. Sometimes our unconscious reinforces our conscious mind and sometimes it overrides it. It gives us intuitions that, research shows, are often remarkably accurate, gives experts abilities that they are unable to explain, and makes most of us avoid doing things we know we should. The American psychologist Jonathan Haidt has a nice metaphor: he sees the conscious mind as a rider on an elephant, a powerful and wilful beast that often decides to disobey the puny being who is trying to steer it. Managers need to understand something of both the rider and the elephant in order to create conditions, tools and processes that lead organisations to learn more from what they do, share individual knowledge more widely, and be more creative.

Learning is a remarkably varied activity: learning a PIN number, a poem, to recognise someone, about the American Civil War, to understand the equations of electrodynamics, to drive a car, to design a building and to manage a project are very different experiences. And whereas most of us can learn to recognise a face without conscious effort, remembering even four random digits requires some conscious attention, learning to drive typically takes 30–40 hours of practice, and it takes over 10 years to qualify as a neurosurgeon (in the UK). Despite this variety, the same memory processes are involved in all of them – and they are central to *using* knowledge, too.

All learning is based ultimately on sensory inputs from vision, hearing, touch, smell, taste and proprioception (awareness of the position and movements of one’s own body). These are stored first in

sensory memory, which is quite capacious but very brief – less than a second for visual memory and only a few seconds for sound memory. Sights and sounds (the only two inputs that concern us here) are processed locally to some extent, and then an interpretation of the parts to which we are paying attention is passed on to working memory. The hearing system, for example, has to disentangle a single stream of incoming sound into the voice on the phone, background chatter in the office, traffic outside and so on. The qualifications ‘interpretation of’ and ‘to which we are paying attention’ are important; we remember only a fraction of what we see and hear, and our memory of an event may be quite different from other people’s, and even quite wrong.

Differences and errors in interpretation are the basis of sensory illusions such as the drawing that sometimes looks like a vase and sometimes like two profile heads facing each other, the Escher staircase that keeps on climbing as we follow it round, only to end up where it started, and the musical tone that goes on rising for ever. Failures of attention are famously illustrated by Harvard professor Dan Simons’ ‘Gorillas in our midst’ experiment, in which about half of the people watching a short video of students playing with a basketball failed to notice a woman in a gorilla suit walking across the scene – even in a variant in which the gorilla stopped half way, turned towards the camera, and beat a tattoo on its chest. Everybody, of course, sees it when it is pointed out.

The visual system constructs its interpretations from numerous ‘snapshots’ in which the eye rests briefly on one point at a time before moving abruptly on to another, and studies of eye movement show that even these depend on what we are paying attention to. People asked variously just to look at a picture, to estimate the economic status of the people in it, or to judge their ages showed quite different patterns of eye movement as they subconsciously searched for different kinds of evidence. It seems that even at the most basic level – while data is still largely sensory, and before significant meaning has been attached to it – the way we perceive the world depends on what makes sense to us, what we expect to see (and hear), and what we choose to pay attention to, almost as much as on what is actually in front of us. No wonder early explorers brought back weird and wonderful accounts of what they had seen, learner drivers find busy towns so confusing, and witnesses can give such different accounts of crimes.

Effects like these are compounded by the characteristics of working memory, with which sensory memory works closely. This is both the next staging post for most of the things we see and hear on their way to long-term memory (some appear to have a more direct route) and the place to which we have to recall existing memories in order to use them in conscious thought. It is where we take the crucial step of

attaching meaning to sensory inputs and encoding them into forms that can be stored in long-term memory, such as words. Information in working memory lasts longer than it does in sensory memory, but not much: about 10–15 seconds. To hold it there longer we have to refresh the memory by repeating it to ourselves. Without rehearsal, the proportion of people who can accurately recall a short, meaningless string of letters slumps from around 90% immediately after seeing them to less than 10% after 15 seconds.

George Miller showed 50 years ago that few people can hold more than seven or eight random independent items (such as random consonants) in working memory, however hard they try, and subsequent research has confirmed that this is about the limit of its capacity.<sup>1</sup> That is extraordinarily small for a channel through which most of our knowledge, memories and conscious thinking have to pass. It would be completely inadequate if memories were like DVDs (just one of which contains nearly 5 *billion* independent items of data), but they aren't. To overcome the capacity limitations of working memory, the brain makes each of the seven or so items it can handle at a time do more work than a byte on a DVD by ruthlessly discarding what it considers irrelevant, encoding information with extraordinary efficiency, and reusing existing knowledge when it can.

Memory of sensory experiences, of course, is only one, limited, kind of knowledge. Neuroscientists distinguish between episodic, procedural and declarative memories, respectively memories of events, motor skills such as driving a car, and verbalisable facts – in common parlance, information – such as names and mathematical methods. These can all be important in professional life (some of the most valuable skills involve combinations of all three) but declarative memory is the most commonly important, and the one that we use most in conscious learning. Research suggests that the keys to this are recognising patterns in what we see, interpreting it in the light of prior knowledge, and connecting it to prior knowledge. A random number generator asked for nine digits is no more unlikely to produce 111222333 than 736129554: both sequences are equally meaningless, but the pattern in the first one – short, simple groups of ones, twos and threes, arranged in a familiar order – makes it much easier to remember. Familiarity can be just as helpful. Many of us use numbers with a personal significance as PINs because they are easier to remember, and it would be much more difficult to learn how to use new software if we did not already understand menus, toolbars, scroll bars and common commands like 'open' and 'save'.

<sup>1</sup> And we can only change its contents about 18 times a second, so our conscious mind can only handle in the order of  $7 \times 18 = 126$  bytes per second of data, a small fraction of the capacity of the slowest modem.

Concepts, language, mathematics, physical laws and rules of thumb are all aspects of these processes at work. They help us to see patterns in observations that would otherwise appear mysterious, to make sense of new experience and knowledge, and to share our understanding. In a famous early experiment on learning a century ago, two groups of children were asked to throw darts at a target underwater. Those in one group were taught about refraction beforehand, so they knew that the actual position of the target would be offset from its apparent position, and that the offset would increase with the depth of water; those in the other group were not. Both groups performed equally at first, gradually learning to adjust their aim to compensate for the water. But when the depth was changed the children who understood refraction were able to adjust their aim more quickly than the others, because they had a mental model of how the water was affecting what they saw. Both as individuals and as a society we build our knowledge and capabilities brick on brick, new on old, using simple ideas to create more complex ones.

Unfortunately, though, understanding an explanation or spotting a meaningful pattern in our own experience is no guarantee that we will remember it in the long term, or even tomorrow. Forming secure long-term memories usually takes practice: we need to recall our new understanding repeatedly over a period of time to develop the necessary neural connections. The more we recall it and link it to other experience and knowledge – by thinking about it, using it, or discussing it with other people, for example – the more securely we remember it, and the more easily it comes to mind when we need it. As Confucius is alleged to have said, ‘I hear and I forget. I see and I remember. I do and I understand.’ The doing is crucial.

## What makes an expert

As people learn more about a field, and gain more experience in it, they become increasingly competent and, eventually, expert. Experts possess vastly more information, examples and mental models than novices, more richly interlinked, and this enables them to *think* differently, too. The difference is perhaps easiest to see in chess players. A novice looks at all the pieces he could move and considers consciously where they could go, what advantage each move might give him, and how his opponent might respond. If he has time and a good enough memory he might go on to consider his options for his next move after that, or even one move further again. Until the game is nearing checkmate, with only a few active pieces left on the board, this approach involves daunting numbers of possible moves and responses, numbers that escalate massively with each additional move considered. The mind boggles at the grandmaster Capablanca, playing a

group of opponents simultaneously, taking two or three seconds to make his move (while everyone else has as long as it takes him to complete a circuit to make theirs), and winning every game. How could he possibly analyse so many possible moves, so quickly and so effectively? The answer is that he didn't. When asked, he is said to have replied: 'I see only one move ahead, but it is always the correct one.' If he was right, chess masters must think very differently from novices. Research suggests that (allowing for a little exaggeration) he was right, and that masters in every field really do think differently. But *how* do they think differently? What make an expert expert?

There is nothing fundamentally wrong with the amateur approach of comprehensively analysing possibilities except the limitations of the human brain, at least in a situation governed by a few formal rules, like chess. Computer chess programs that work in more or less that way can be quite strong players: fast processing and perfect memory make it possible for them to look much further forward than humans, and though their play lacks flair, it is effective. But tests show that grandmasters do not think faster, or have better memory in general, than other people, and yet they can still beat computers most of the time. Research has shown that they often do analyse possible moves, but no more of them than moderately competent players. However, they concentrate their analysis on the most promising ones – and sometimes (in Capablanca's case, usually) they can see the best one immediately.<sup>2</sup>

Grandmasters' years of thoughtful practice and study give them a hugely greater repertoire of remembered positions, moves and strategies than the average player, and that enables them to short-circuit analysis by recognising patterns and recalling ready-made solutions. They make extensive use of their long-term memory while weaker players have to make do with working memory. Experts in other fields do the same: when presented with a task, their experience and deep understanding enable them to see the features that matter, ignore irrelevancies on which the less expert tend to get hung up, home in on the factors that are likely to lead to a solution, and make progress quickly. They have the kind of intuition that Nobel prize-winning psychologist Herbert Simon called 'analyses frozen into habit and into the capacity for rapid response'.

Recent research on mice brains by Joe Tsien and his colleagues at the University of Boston suggests that this ability is inherent in the way that their – and our – brains form memories. Using sophisticated experimental techniques, instrumentation and mathematical analyses

---

<sup>2</sup> Quite a contrast to IBM's Deep Blue, which needed to evaluate 200 million moves a second to become the first chess-playing computer to beat a reigning world champion in 1997.