



Part One
Foundations

COPYRIGHTED MATERIAL

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.¹

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

¹ 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.