

Introduction

You can see a high-performance factory or office, but it just doesn't spread.
I don't know why.

William Buehler, Senior Vice-President at Xerox

Mr Buehler's observation (Jacob, 1992) is a contemporary expression of an old puzzle – one that hasn't yet been resolved. Indeed, organizations often do not have to look too far to find best practices. In many cases, they find stellar performance in their own backyard. It seems sensible to expect that in-house examples will diffuse to other units of the organization, once uncovered. Peers will imitate and management will 'suggest'. Yet, evidence shows otherwise. Best practices do not readily spread within firms.

One reason for this might be that companies simply do not attempt to spread best practices. Incentives to search for better practices inside the firm may be limited because, in the absence of compelling evidence, comparable operations are expected to have equivalent performance. For example, similar semiconductor plants – with comparable equipment, personnel and technology – are mostly expected to yield similar productivity and quality. Thus, organizations with multi-plant or multi-office structures see themselves as homogeneous collections of similar units and, understandably, turn outside for inspiration.

Even when internal results clearly stand out, limited understanding of the underlying processes and lack of adequate measures make it laborious, uncertain and generally unfruitful to advocate the transfer of those practices to other sub-units of the organization. Without timely, detailed and comparable measures of operational performance (Jacob, 1992; Kaplan, 1990), the economic magnitude of gaps in operational performance can easily be dismissed as exaggerated, manipulative, or as the inescapable consequence of structural, unconquerable factors (Chew *et al.*, 1990; Hayes and Clark, 1985). Thus, when the tangible trauma of change is pitted against speculative benefits, efforts to identify and replicate superior practices within firms are frequently relegated to the category of 'important but not urgent'.

This reality has been slowly changing. Fact-based management techniques brought by total quality programmes (Crosby, 1984; Ishikawa, 1985; Juran, 1988), benchmarking initiatives (Camp, 1989) and re-engineering

(Hammer and Champy, 1993) have improved dramatically the understanding of internal operations and the availability of fine-grained measures of operational performance. Timely collection, dissemination and use of the information generated from these measures is now possible with reasonable effort, thanks to advances in information systems, e.g. intranets, data warehouses, decision support tools, ERP systems and group-ware. Gaps of 200 per cent or more in the performance of comparable units¹ – gaps worth several million dollars (Chew *et al.*, 1990) – are frequently found. The prospect of financial gains of that magnitude naturally triggers efforts to narrow these performance differentials.

For these reasons, a large number of organizations are now attempting to transfer best practices, to close internal performance gaps, to stop re-inventing the wheel and to eliminate deficiencies in performance. The rise of the knowledge economy has helped organizations recognize that knowledge assets are rapidly becoming their most precious source of competitive advantage, and that learning to better manage those assets has become a competitive necessity. Accordingly, it is increasingly common to find executive positions, such as Chief Knowledge Officer or Chief Learning Officer, that have the explicit mandate to transfer existing knowledge to other parts of the organization.

Yet, even though more attention is being directed to best practices, these remain stubbornly immobile. In a survey of 431 US and European organizations conducted in 1997 by Ernst & Young, only 14 per cent of the respondents judged satisfactory the performance of their organization in transferring existing knowledge internally. The remaining 86 per cent found it lacking (Ruggles, 1998). Another survey of 79 subsidiary presidents and their immediate superiors of three global Fortune 500 corporations found big gaps between expectations, perceptions and reality. Whereas the parent company expected 95 per cent of the subsidiaries to be actively sharing knowledge and perceived that about 89 per cent were actually doing so, in reality only 62 per cent were actively engaged in knowledge-sharing activities (Gupta and Govindarajan, 2000). Best practices are unlikely to spread if companies do not try to spread them. However, even when they do try to spread them, best practices spread less than expected, because transferring them effectively is often found to be far more difficult than expected. Transfers of practices within the firm tend to be ‘sticky’.

This book

Why don't best practices spread? Any progress that could be made in understanding and unlocking the puzzle will have implications for strategic management, organizational theory and ultimately for society at large. That is because the notion of stickiness is related to fundamental questions such as why and when a firm may be superior to a market in creating and

transferring knowledge, how organizations learn to derive competitive advantage from their knowledge resources and the general societal concern of how to better utilize existing knowledge assets. Thus, for example, understanding stickiness could help us better appreciate the workings of organizational flexibility, the potential value of acquisitions, the chances for success of strategic alliances, technology partnerships and technology transfer agreements and, more broadly, how organizations leverage their knowledge.

It is thus not surprising that the best-practices puzzle, articulated more than three decades ago by Dick Walton from the Harvard Business School, emerged as one of the most important managerial challenges of the late 1990s and remains high in the priority list of the new millennium (see, for example, Cairncross, 2000; Slywotzky and Morrison, 2000; Stewart, 2000).

This book is a close and careful look at the best-practices puzzle. You will find an in-depth look at transfers of best practices, a detailed exploration of the nature of the difficulties that might be encountered, of the factors that may underlie those difficulties and of some of the possible explanations for the puzzle that those factors suggest. Both qualitative and quantitative methods were used to understand the puzzle and look for possible ways to explain its persistence.

The persistence of the best-practice puzzle is in itself puzzling because the observation that it is difficult to transfer best practices is hardly a new discovery. Indeed, implementing internal transfers of best practice – of a superior technology or of a better way to organize work – have long been recognized as an important managerial challenge. As early as 1913, Ford transformed its entire operation from craft to assembly-line production (Hounshell, 1984). Toyota diffused the *kanban* system throughout the company (Ohno, 1978). Yet, in both cases, competitors encountered great difficulties in imitating these practices. Outside the automobile industry, there were public attempts to replicate exceptional manufacturing practices from DEC's Enfield factory, GE's Bromount factory and Westinghouse's College Station factory (Ulrich and Lake, 1990).²

The puzzle may persist because factors that could be grouped under the rubric of motivational barriers are typically the only ones blamed for the lack of diffusion of practices. Difficulties to the transfer of best practices within firms are traditionally ascribed to interdivisional jealousy, lack of incentives, lack of confidence, insufficient priority, lack of buy-in, a heavy inclination to re-invent the wheel or to plough twice the same fields, refusal of recipients to do exactly what they are told, resistance to change, lack of commitment, turf protection and many other manifestations of what seem to be part of the popular definition of the Not-Invented-Here or NIH syndrome.

Researchers who have looked at the phenomenon seem to agree. For example, Michael Porter notes that 'the mere hope that one business unit might learn something useful from another is frequently a hope not realized' (1985: 352). He explains that '[b]usiness units acting independently

simply do not have the same incentives to propose and advocate strategies based on interrelationship as do higher level managers with a broader perspective'. He blames both the recipient, who can 'rarely be expected to seek out know-how elsewhere in the firm', and also the source, who 'will have little incentive to transfer [its know-how], particularly if it involves the time of some of their best people or involves proprietary technology that might leak out' (1985:368).

When difficulties are primarily pinned down on motivational factors, adequate incentives appear indispensable. For example, Porter argues that '[u]nless the motivation system reflects ... differences [in perspective], it will be extremely difficult to get business units to agree to pursue an interrelationship and to work together to implement it successfully. Instead they become embroiled in fruitless negotiations over the allocation of shared costs or over procedures for sharing revenue' (1985:386). In the same vein, Goold *et al.* (1994) note that enlightened, self-interested business unit managers will exert their implicit veto rights on opportunities for knowledge sharing that they personally find unattractive. Thus, overcoming difficulty is tantamount to convincing those business unit managers.

Approaching transfers of knowledge with such an exclusive focus on incentives immediately directs attention to transfer-related benefits that are or appear asymmetric, and to corporate incentive schemes that, by not offering any specific incentive to transfer, indirectly penalize managers who incur costs in supporting them. Corporate management is often reluctant to modify incentive systems because it fears that treating business units differently, or creating idiosyncratic measures of performance for each unit, will vastly complicate the management of the organization. That is because units that are sources of best practice might be able to excuse poor performance by citing their efforts to aid others. Thus, corporate management, rather than tinker with the organization's incentive system, prefers to leave the situation as it is. Maybe that's why the puzzle has persisted for such a long time.

In my quest to explore the best-practices puzzle, I have naturally considered the impact of motivational barriers, but did not stop there. Through a careful and extensive review of related literature and evidence on knowledge transfer and on how corporations use their knowledge assets, I discovered another set of reasons, besides incentives, that may explain why knowledge might not transfer.

I call this alternative set of reasons *knowledge barriers*. Examples of knowledge barriers are the recipient's level of knowledge prior to the transfer, how well the transferred practice is understood within the organization, the recipient's ability to unlearn, i.e. shed prior practices, and the pre-existing social ties between the source and the recipient of knowledge. These factors are qualitatively different from motivational barriers, such as the motivation of a source to share knowledge or to support the transfer of that knowledge and the motivation of a recipient to absorb and institutionalize external knowledge.

When motivational and knowledge barriers are both taken into consideration, a different picture emerges. Indeed, I've found that knowledge barriers could overshadow motivation barriers to the transfer of best practices within the firm, a discovery that has far-reaching implications for those who grapple daily with the best-practices puzzle and seek effective ways to enhance best-practice sharing and the use of existing knowledge within their organization. That is because it seems that it might be possible to design several alternative ways to enable knowledge sharing that do not require any modification to the incentive system and thus could be implemented within the existing organizational structure.

That basic picture can be elaborated further by taking into account the different stages that typify the evolution of a transfer. When I did that, I discovered that the relative importance of each type of barrier changed with the stage of the transfer, again highlighting non-obvious dynamics (e.g. that a motivated recipient could actually cause difficulties to transfer knowledge) and areas for future research. A more nuanced picture of the transfer suggests opportunities for sophisticated managerial interventions.

In sum, this book is a careful and detailed exploration of the best-practices puzzle. In the first part of the book, I discuss how I have approached the study of the puzzle, i.e. I define stickiness, describe the kinds of stickiness that one might consider and what kind of barriers one might expect. I then show both qualitative and quantitative evidence of stickiness and of its predictors, and discuss the implications that these findings may have for both research and practice. To help the reader further interpret the evidence I provide, I have included abundant detail of the methods that I have used to conduct this research.

I embarked in this quest because I was really intrigued by the persistence of the best-practices puzzle, by the seemingly limited effectiveness of conventional remedies and by what that implied about an organization's true ability to leverage existing knowledge. I attempted to go beyond just trying to provide one more key to unlock the one and only known gate to the effective transfer of best practices. I sought alternative gates. I believe that I have found some non-obvious ones, a discovery that in turn opens exciting alternative ways to leverage knowledge within the firm.

Organization of the book

Because the topic of this book speaks to a variety of practical and academic concerns, I structured the book so that it could be conveniently accessed by different readers, both by those with an academic orientation and by those with a practical one. For this reason, the main topics, ideas, findings and conclusions are covered in a relatively brief and accessible format in the main body of the book, which is followed, in technical appendices, with abundant detail about methods.

Thus a reader with practical preoccupations may read Chapter 2 for definitions, and then turn immediately to Chapter 7 for an overview of the findings, referring back to Chapters 4 and 5 to fill in details, to Chapter 6 for concrete examples and to Chapter 8 to read about the practical implications of the findings. This will give a practically minded reader enough familiarity with the topic to begin to relate the contents of the book to a specific situation, or perhaps to begin to sketch possible alternative courses of action.

Practical implications could be made much more detailed and specific by collecting and analysing information from a given situation. Such an in-depth quest will require a more careful reading of Chapter 4, which provides a conceptual discussion of the different barriers to knowledge transfer. The technical appendices contain tools that can be used to measure those barriers. In particular, the appendices contain an elaborate questionnaire with measures for each of the barriers, as well as for many other aspects of the transfer that could help paint a rather comprehensive picture of the transfer situation. Information collected with those tools could be then used to find out what happened in a particular instance or to inform an effort to identify and preempt difficulties.

The technical appendices, besides data collection tools, include a detailed exposition of the research methods as well as additional detail on the statistical findings. Those interested in researching knowledge within organizations may find in the appendices a description of special challenges that emerge in this kind of research and of how some of those challenges were met in this particular study.

Several other aspects of the book, besides the technical appendices, should be particularly appealing to graduate students. For example, Chapter 3 positions knowledge transfer within the concerns of strategic management. Further, the text offers a rather comprehensive review of relevant literature, especially in Chapters 4 and 5. In addition, for those interested in stickiness research, Chapter 8 offers a number of research implications, and sketches possible avenues for further research.

Finally, teachers both at the graduate and the undergraduate level may find that the book provides a general introduction to knowledge transfer for their students. Furthermore, parts of the book could prove to be useful additions to courses that span a variety of knowledge-related topics, such as knowledge management, organizational learning, benchmarking and the sharing of best practice.

In the end, I realize that each reader's needs are in some sense unique and could span a broad range of interests. Thus, rather than speculating further about different strategies for reading the book, I give below a brief description of the contents of each chapter to help the reader tailor his or her approach to specific concerns, angles, questions and interests.

In *Chapter 2*, I argue that the transfer of knowledge within the firm can be difficult. I then define the notion of stickiness as the difficulty to transfer knowledge and discuss how one could detect the existence of stickiness.

Chapter 3 positions the transfer of knowledge within the field of strategic management. The thesis of this chapter is that stickiness reflects the presence of internal factors that impede the realization of competitive advantage. It is claimed that stickiness hinders the appropriation of rents from existing knowledge assets. This in turn suggests that factors that cause stickiness act as internal barriers to rent appropriation.

Chapter 4 offers a typology of predictors of stickiness. Stickiness can be predicted by analysing properties of the transfer. In general, the unfolding of the transfer depends to some extent on the disposition and ability of the source and recipient, on the strength of the tie between them and on the characteristics of the object that is being re-created. The features of the organizational context where re-creation occurs are important as well. The impact of these factors is analysed in this chapter.

Chapter 5 offers a typology of stickiness. Four transfer phases are synthesized from the literature on knowledge transfer: initiation, implementation, ramp-up and integration. Each of these four phases can be difficult in a different way. Accordingly, I distinguish between four types of stickiness: initiation stickiness, implementation stickiness, ramp-up stickiness and integration stickiness. The nature of difficulty at each stage and possible predictors are discussed.

Chapter 6 illustrates the four types of stickiness – initiation, implementation, ramp-up and integration – with findings from in-depth fieldwork in three different companies. Initiation stickiness is illustrated with Rank Xerox's difficulties to initiate transfers between its European subsidiaries. Implementation and ramp-up stickiness are illustrated with Banc One's difficulties experienced when converting acquired banks. Integration stickiness is illustrated by Centel's difficulties to sustain in each of its divisions a 'best' practice that had already been effectively implemented.

Chapter 7 presents the results of statistical analysis aimed at identifying which were the best predictors of difficulty for each stage of the transfer and overall for the transfers of this study.

Chapter 8 suggests some implications of this study for further research on stickiness and for the practice of knowledge transfer. The chapter discusses the contributions to extant research on stickiness, which include a typology of stickiness as well as alternative ways to measure stickiness. The discussion of implications for research ends with suggestions for more specialized research on stickiness and on its antecedents. The development of implications for practice includes speculation about possible actions that could be taken to facilitate each stage of the transfer. Next, practical questions that typically arise during the initiation of knowledge transfer are addressed using data collected in this study. This includes the question of who should be the first recipient of an internal best-practice transfer when there are several viable candidates to choose from. It also includes the examination of how senior management interventions could affect stickiness.

Chapter 9 provides a summary of the book, the underlying study, its findings and their significance for research. The chapter concludes by arguing that we now have new clues to rethink prevailing wisdom about why best practices may not spread. Conventional wisdom blames incentives, i.e. motivational barriers, almost exclusively. The findings, however, point to knowledge-related barriers, not just to motivation-related barriers, as an important culprit.

The book also contains five technical appendices with details of the study. These appendices contain the research design, the questionnaire used for Phase I, an example of a cover letter used inside the companies, a partial list of the practices studied and a full version of the questionnaire used in Phase II of the study.

In the next chapter, I expand on the best-practices puzzle and introduce the notion of stickiness.

Notes

1. Besides the published references, I've found up to 10:1 gaps in performance in otherwise comparable units, and gaps of 2:1 rather frequently. Personal communication with Robert Camp, a widely known benchmarking specialist from Xerox, confirmed that gaps of 200–300 per cent are a typical finding in internal benchmarking efforts.

2. Likewise the socio-technical literature describes numerous attempts to replicate novel ways of organizing work internally (Walton, 1975).