

WORKING WITH TACIT KNOWLEDGE

Joseph A. Horvath, Ph.D.

IBM Institute for Knowledge Management

WORKING WITH TACIT KNOWLEDGE

Introduction

Transforming changes in the business environment have brought the issue of organizational knowledge into focus for researchers, consultants, and managers. It has become widely accepted that businesses compete on the basis of knowledge and, indeed, that knowledge may be their only sustainable source of competitive advantage. It has become likewise accepted that knowledge or “know how” represents an invisible but substantial component of a firm’s market value and, thus, an asset to be measured and managed—just as assets are measured and managed. This article explores some of the conceptual and practical issues that companies must confront when they begin to look deeply at what it takes to manage knowledge. A premise of the article, amply supported by research, is that some of the most valuable knowledge within a firm is essentially hidden or tacit—residing not in documents or databases but in the experience and skill of human beings. To talk of “managing” such knowledge is somewhat misleading; it is perhaps more accurate to speak of “working with” tacit knowledge. The goal of this article is to introduce readers to some of the concepts, methods, and practices required to work effectively with tacit knowledge in business settings.

Philosophical debates about the “true” nature of knowledge have a long and interesting history but are, in business settings, mostly beside the point. To work with knowledge a simple, working definition is required. One basis for definition is to recognize that people add value to information and, in so doing, create knowledge. People combine information with other information to form new and unique combinations. They refine information for specific uses or generalize it for broader application. They evaluate or “warrant” information to establish its truth or usefulness. Occasionally, they reformulate or reframe information to yield new insights. In these and other ways, human beings lend context, meaning, and purpose to information and thus move it along a continuum toward what we commonly call knowledge. Thus, a working definition of knowledge—suitable for use in a business setting—might be *information with significant human value added*.

For example, medicinal drugs reach the market accompanied by printed package inserts that describe indications and contraindications, dosage instructions, and possible side-effects. From the perspective of the scientists and regulators who compose these inserts, their contents represent *knowledge*. That is, they are the result of a value-adding process whereby a vast array of information (e.g., on dose response, safety, clinical efficacy) is organized, validated, and applied to build a usage profile for the drug. From the point of view of the physician, however, the package insert represents *information*. That is, it represents the starting point in a value-adding process in which the physician combines information from the insert with other

information (e.g., on symptom patterns, outcomes for different patient subtypes, cost considerations) as she learns from experience about when and how to use the drug.

As these examples suggest, we can distinguish knowledge from information, but only with reference to a particular intended use. Philosophers may define knowledge in *structural* terms (i.e., in terms of its relation to other concepts) but, in business settings, it makes more sense to define knowledge in *functional* terms (i.e., in terms of its use). This means that what is knowledge in one context and for one purpose may be information in another. There is nothing particularly mystical about such a definition. Consider the “value chain” whereby raw materials and human labor are progressively converted into goods and services. Any link in that chain may be considered either a “customer” (for the purpose of describing its upstream role) or a “supplier” (for purposes of describing its downstream activities). The conversion of information to knowledge occurs along just such a value chain and the terms “knowledge” and “information” are similarly dependent upon the context in which they are used.

The Tacit Dimension

Although disagreement about the definition of knowledge persists, there is much greater agreement about the type or form of knowledge that yields the greatest (potential) business value. Thus, although many of the early forays into knowledge management centered on the use of information technology to broaden access to documents and databases, there is by now widespread agreement that much of the highest-value knowledge within organizations is uncodified¹. Quite simply, lots of valuable knowledge “falls through the cracks” within business organizations, never finding its way into databases, process diagrams, or corporate libraries. As a consequence, much of what the firm “knows” remains unknown or inaccessible to those who need it. Such knowledge is present within the organization, but it remains hidden, unspoken, tacit. In business organizations, this hidden or *tacit knowledge* takes one of two forms: 1) knowledge embodied in people and social networks, 2) knowledge embedded in the processes and products that people create.

Embodied Knowledge

People know more than they can say. That is, they are guided by knowledge that they do not know they have or that they find difficult to articulate.² An expert drill-press operator, for example, can accurately predict when a drill bit will reach the end of its effective life. Because she possesses this knowledge, the expert operator uses fewer drill bits and ruins less material due to breakage than does the novice operator. Simply asking the expert operator what she knows, however, will yield disappointing results. She’ll tell you that it’s simply a “feeling” that she gets from the noise, vibration, and the smell that the drill press gives off. We may say that such knowledge is “embodied” in that it cannot easily be detached from the knower.

The practices or “ways of the shop” that develop unnoticed over time in an organization also represent tacit knowledge. These may not reside in a single person’s head but, rather, may be distributed across a group of

people and we may say that such knowledge is “embodied” in groups, teams, or communities.³ In continuous-process manufacturing, for example, operator/controllers develop an unspoken understanding for each other’s working styles and will take these into consideration at shift changes. Groups that have worked together for a long time tend to “reboot” a plant less frequently because they can understand and accommodate themselves to control settings established by the prior shift. These groups achieve better yields and lower energy consumption by virtue of the mutual understanding that has developed over time. This form of embodied knowledge—the knowledge that resides in groups, teams, or communities—is a key source of under-leveraged know-how in most organizations.

Embedded Knowledge

Beyond the knowledge that resides in people, we must also consider knowledge that resides in the things that people have produced. Knowledge may be embedded in products or product prototypes, in processes, as well as in documents (i.e., in the sense that the background or context necessary to understand a document may not be explicitly given). Although these may be poor proxies for embodied human knowledge, they are important mediators of knowledge transfer in all business settings. For example, Dorothy Leonard⁴ has shown that equipment or product prototypes serve as important “boundary objects” in the problem solving of cross-functional teams. Such prototypes allow individuals and groups to externalize their tacit knowledge—in the form of ideas, assumptions, and value judgments—so that these can be reconciled and incorporated into subsequent versions of the prototype.

Processes are also sources of embedded knowledge and can serve as boundary objects in knowledge transfer. In pharmaceutical process development, for example, simpler and more manufacturable “synthetic routes” are sought as a lead compound progresses toward clinical trials. These synthetic routes (series of chemical reactions which add or subtract atoms from a molecule) are the subject of much “learning while doing” and often take into account such unscientific factors as union rules dictating who can touch what in the pilot plant.⁵ Although final process specifications will be reported to regulatory agencies and are available to other developers, the story behind those specifications will likely be much less accessible—even though it might help reduce process-development lead times on future projects. In general, we may say that business processes only implicitly reflect the learning that went into their development and that this form of tacit knowledge represents an untapped resource in many companies.

In summary, the term “tacit knowledge” has come to refer quite broadly to the know-how that is hidden or implicit in organizations. Formal knowledge in official documents and databases represents a small fraction of what an organization “knows.” To extract more business value from their knowledge, organizations must find ways to penetrate the sticky, opaque character of knowledge—its tacitness. As I argue in the section to follow, they may expect to be rewarded for these efforts.

The Business Value of Tacit Knowledge

There are several reasons why tacit knowledge merits deliberate management in business settings. A great deal has already been written on this topic but, for brevity, we may identify several broad categories of business benefit. In the section that follows, I describe several notable examples of organizations that have begun to realize these benefits.

Innovation

Tacit knowledge is strongly implicated in organizational innovation. People develop and use tacit knowledge before they are able to formalize or codify it. Thus, the leading edge of the firm's learning (and a source of its future innovations) is often to be found in the tacit knowledge of its people. In a widely cited example of how tacit knowledge can spark innovation, Matsushita sent a senior engineer to apprentice with a master chef at a top Tokyo hotel. The tacit knowledge acquired in this apprenticeship—knowledge of the look, feel, and taste of bread dough—was subsequently applied to the design of an innovative “kneading” mechanism for an automatic bread-making machine⁶. More often, as Nonaka and Takeuchi have shown, Japanese firms work with tacit knowledge by leaving room for it in their management structures and communications. For example, firms such as Matsushita and Honda employ figurative or metaphorical language (e.g., “rich and easy”) to express product concepts early in their development. In contrast to rigid design specifications, these somewhat fuzzy product concepts admit of varied interpretations and, thus, stimulate the articulation and development of knowledge that might otherwise remain tacit. Although innovation remains a poorly understood subject, the best available research suggests that tacit knowledge, properly mobilized within an organization, drives a “virtuous spiral” of further knowledge creation and innovation.

Best Practices

Attention to tacit knowledge can enable firms to identify and transfer best practices more effectively.⁷ People develop tacit knowledge as they solve real problems in pursuit of real goals. This means that tacit knowledge, when compared with explicit knowledge or information, tends to reflect more closely the reality of how work actually gets done (i.e., work “practices” rather than business “processes”). Often, what is “best” about a practice fails to show up on process maps or equipment specs.⁸ A study by Wenger of claims processors in a health-insurance company shows that effective practice—at both the individual and group level—is bound up in the acquisition and exercise of tacit knowledge. Claims processors develop tacit knowledge of patterns that recur in the endless flow of claim forms (e.g., which claims look problematic and so should be put aside for later in the workday). They learn shortcuts and “work arounds” that allow them to meet their production quotas without substantively compromising quality. Perhaps most importantly, they participate in a “community of practice” in which tacit knowledge is distributed across individuals such that no single individual needs to know everything. In short, excellent business practices cannot be effectively

transferred unless they are well understood. And effective practices cannot typically be understood without reference to the tacit knowledge of the people who do the work.

Imitation

Tacit knowledge can help firms to resist imitation by competitors.⁹ Because it is embodied in people and embedded in the things they create, tacit knowledge tends to be “sticky”—to resist transfer to new groups and settings. Although this stickiness makes the mobilization of tacit knowledge particularly challenging, it makes its appropriation by competitors even more challenging.¹⁰ Thus, firms that work effectively with tacit knowledge can expect to increase both their ability to innovate and their ability to extract innovation rents in the marketplace. At Chaparral Steel, for example, management offers guided tours of its plant to anyone—including managers from competing firms. This open door policy reflects a well-founded belief that Chaparral’s sources of advantage are not easily replicable. They are not replicable because they reside largely in the tacit knowledge of the firm (both embodied and embedded) and in the set of mechanisms that management has put into place to support the creation, sharing, and application of that knowledge. Because much of its knowledge resides in people (and because Chaparral is better at leveraging that knowledge than anyone else) it has more to gain than to lose by opening its doors¹¹.

Core Competencies

Finally, a consideration of tacit knowledge can illuminate the emerging core competencies of the firm. Tacit knowledge represents the unique value added by the people who generate it. It emerges from their particular situations, skills, and experiences and, in aggregate, reflects the history and circumstances of the firm. In this sense, tacit knowledge needs to be considered in the evaluation of the firm’s core capabilities—those “best in world” capabilities with the potential to distinguish the firm from its competitors. By getting a handle on its tacit knowledge assets, a firm can better understand its competitive position and can more effectively select and shape the markets in which it competes. An example of a firm (Xerox) deriving value from tacit knowledge in this way is provided in the section that follows.

Working with Tacit Knowledge: Notable Practices

Firms don’t work with tacit knowledge for its own sake, nor even to realize the somewhat abstract benefits described above. Rather, most successful efforts to date have used the recognition and leveraging of tacit knowledge to add power and depth to existing projects for which a business case has already been developed. Three projects—individually notable and collectively representative—are summarized in this section. They range from strategic planning to process improvement to culture change and offer the reader some sense of how, and how intensively, some leading organizations are working with tacit knowledge.

Improving the Submissions Process for New Drug Approval

Pharmaceutical companies live and die by their ability to speed drugs to market. The patenting of a new chemical entity typically precedes its final approval as a drug by several years and the remaining period of patent-protected sales is critical to overall profitability. Executives in this industry are known to value time-to-market at one million dollars per day and so are acutely interested in anything that will reduce product cycle time. In the area of regulatory submissions, electronic document management systems (EDMS) have been widely and successfully deployed to this end. A process redesign effort at one global pharmaceutical company, however, shows the potential benefits of moving beyond a technology-centered approach to capture and leverage the tacit knowledge that is present but fragmented within a firm.

The objective of the process redesign was to achieve a significant reduction in one segment of product cycle time—the period from the last patient visit in clinical trials to the submission of the New Drug Approval (NDA) dossier. A study of the current state indicated that this filing time averaged 35 to 40 weeks, reflecting in part the high level of rework involved in the submissions process. The study also indicated that the firm's current capabilities were insufficient to support simultaneous filings with multiple regulatory bodies—an impediment in the global marketplace. Thus, although it was focused explicitly on the mobilization of tacit knowledge, the redesign of this company's submissions process was driven by business issues, not by knowledge concepts.

The submissions redesign began with the identification of tacit knowledge, resident within the firm, that was relevant to the development of effective submissions processes. The study of the current state had indicated that tremendous knowledge was indeed present within the company but that this knowledge was embodied (i.e., in a few experienced individuals) and embedded (e.g., in practices that successful teams had converged on as they struggled to bring their particular drugs to market). A first step in leveraging this knowledge was, of course, to assemble a cross-disciplinary group composed of individuals with submissions experience and, thus, relevant tacit knowledge. This group worked to develop guidelines for planning, front-loading, and tracking the documentation activities that precede regulatory submission—bringing their collective experience to bear on issues that, until recently, had been the province of isolated experts or even trial-and-error.

Having mobilized the tacit knowledge of its members, the redesign group used existing technology and new policies to leverage that knowledge more effectively. Documentation guidelines were designed for easy access and readability via the company's intranet—part of a broader effort to provide scientists with knowledge “on demand.” In addition, a debriefing process was instituted for all global project teams and roles and responsibilities associated with knowledge capture were specified. Debriefing sessions were mandated following key project milestones (e.g., filing, regulatory response) and the results of these debriefings triggered evaluation and modification of the documentation guidelines as appropriate. The

integration of expert knowledge with the debriefing process distinguishes this submissions re-design as a fully-fashioned knowledge project. It closes the “learning loop” on regulatory submissions by ensuring that the knowledge of the re-design group will be augmented and refined going forward.

The business benefit delivered by the submissions redesign was clear and compelling. Filing time was reduced from 35 to 18 weeks. By even the most conservative valuation of time-to-market, the company achieved a very high return on their investment in leveraging tacit knowledge. As the project leader put it, “Tacit knowledge has been made explicit, and our processes have been redesigned to reflect that knowledge.”

Capturing “Lessons Learned” in the U.S. Army

Many of the “business” issues that have driven the U.S. Army’s foray into knowledge management will be familiar to those who manage for-profit enterprises. These issues include how to adapt to a rapidly changing competitive environment, how to integrate new technologies into operations, and how to do more with less in an era of reduced headcount. Today, the U.S. Army faces new and very different missions such as those exemplified by police actions in Somalia, Haiti, and Bosnia. The Army’s forces are largely task-organized and faces many of the same command and coordination challenges as civilian organizations. The Army employs a constantly changing array of advanced technologies, with all of the training and human performance challenges that this entails. Finally, the Army bears the historical legacy of the Vietnam war in which it learned the terrible human cost of not learning from its mistakes.

One way in which the Army has addressed these issues has been to systematically improve its ability to capture, transfer, and reuse knowledge. These efforts have been the subject of both research and media attention, but we may fairly summarize their approach as follows. First, the Army *institutionalizes* knowledge capture by building simple but proven practices into its everyday operations. These practices, embodied in the “After Action Review”, are consistently and rigorously applied throughout the Army to ensure that mistakes will be uncovered and successes reflected upon. Second, the Army *amplifies* captured knowledge by creating a limited number of knowledge-management roles which serve to move captured knowledge beyond the local level. These dedicated knowledge-management roles are concentrated at the Center for Army Lessons Learned at Fort Leavenworth, an historical center of officer education and doctrinal development.

The heart of the Army’s knowledge-management efforts is the “After Action Review” (AAR) process. After every mission, project, or exercise, participants come together to address three fundamental questions:

1. What did we set out to do?
2. What actually happened?

3. How do we account for the difference?

The mission or project is broken down (often by task or function) into manageable pieces, each of which becomes the focus of a separate AAR. A discussion leader is designated to conduct a group discussion that is structured around the three questions. Active and candid involvement is expected of all participants, with the understanding that nothing said within the AAR session will influence performance evaluations or otherwise be held against the speaker. Candid speech is protected in the AAR context and withholding information is widely viewed as a career-limiting move. To its great credit, the Army has succeeded in making open and engaged participation in AARs a governing value within the Army.

Participants in an AAR typically hold a great deal of knowledge in common. This knowledge ranges from specific, disciplinary expertise to an understanding of each other's backgrounds and working styles, to a ready familiarity with the problem or "terrain" being discussed. This high degree of shared context eases communication within AARs and makes them powerful settings for the articulation and transfer of experience-based, tacit knowledge. In AARs, as in many close working relationships, people tend to finish each other's sentences and can penetrate deeply into problems without benefit of a long, "level setting" exposition. Whether conducted in a Pentagon conference room or in the back of an armored truck, the After Action Review is how the Army goes about leveraging the tacit knowledge of its people.

Commanders use the output of AAR sessions for a variety of purposes. These purposes include modifying standard operating procedures, promoting knowledge exchange among the units under their command, instructing and developing their subordinates, and bringing their own mistakes and weaknesses to light. What AAR sessions are less useful for, is moving "lessons learned" beyond the local level—amplifying them for consideration by a wider circle of practitioners. Such consideration is necessary not simply to broaden access to local knowledge but also to allow that knowledge to be qualified or "warranted" before it is put into practice widely. Like civilian organizations, the Army must balance the control that traditional approaches to warranting (i.e., passing things up and down the hierarchy for approval) confer, against the advantages in speed and flexibility that more lateral, decentralized approaches confer.

In order to drive more effective knowledge management on an "enterprise" level, the Army has established a core team of experts, located at the Center for Army Lessons Learned (CALL). The members of CALL do not manage the Army's knowledge, however. Rather, they facilitate and coordinate the activities of an extended team—assembled on-demand from the Army's vast network of practitioners and subject-matter experts. The extended team is sent into the field to observe, interpret, and disseminate lessons learned in important operations. Just as the Army "spins up" an expeditionary force to address military threats anywhere in the world, so it spins up a learning force to address important learning opportunities.

This learning force, trained and coordinated by CALL, has three objectives. First, they endeavor to be there when interesting things happen. This requires some “pre-processing” on the part of the learning force. For example, subject-matter experts who are tasked to support a CALL engagement must develop “issue lists” that identify where the Army’s knowledge may be thin or out of date. Second, the learning force works hard to make sense of events—to put them into context, to test their understanding and to warrant them through circulation to a widening network of interested and knowledgeable parties within the Army. Finally, the learning force seeks to put the lessons they have distilled into the hands of those who need them. This transfer of knowledge to its end users ranges from the very local (e.g., *ad hoc* briefings delivered “on the ground” during an engagement) to the very global (e.g., reports, practitioner bulletins, and multimedia training materials available on an Army-wide basis).

The work of the Center for Army Lessons Learned (and that of the extended network they draw upon) has resulted in several highly visible “wins” for the U.S. Army. In the deployment of peace-keeping forces to Haiti, for example, the training scenarios developed by CALL (who accompanied the first U.S. troops onto the island) were instrumental in preparing replacement troops to succeed in a complex, potentially explosive situation and thus in minimizing human casualties—truly the ultimate “cost” consideration. Highly visible successes such as these have also helped the Army to build a reputation as an exemplary learning organization, with very favorable profiles in outlets such as *The Wall Street Journal*, and *The Harvard Business Review*.

Auditing “Specific Knowledge” at Xerox

Like all technology companies, Xerox Corporation operates in a highly fluid, highly competitive environment. New and potentially “competence-destroying” technologies wash up on its shores each year. Every marginal dollar of profit in its core businesses is fiercely contested. Thus, despite its early technology leadership and high name recognition in the market for photocopying equipment, strategic planning at Xerox cannot be viewed simply as a “war of position” (i.e., selecting markets and erecting barriers to entry) but must instead be understood as a “war of movement” (i.e., anticipating possible futures and building the capabilities needed to get there first).¹²

Consistent with this resource-based view, Xerox sought to support its strategic planning process with an assessment of the current and future competitive status of its core capabilities. The capability to produce “marks on paper” is one example of such a capability and Xerox wanted to understand the nature and status of the knowledge assets that made up this capability--wherever they might reside within the company. Beyond these strategic considerations, Xerox wanted to derive more value from its intellectual property by broadening its definition to include “specific knowledge” (tacit knowledge of both the embedded and embodied variety) and by implementing more deliberate management of these assets.

In describing Xerox's approach to knowledge auditing, we will use the paper-marking competency as an example. This competency was decomposed to form 28 distinct specific-knowledge categories. These categories represented domains of tacit knowledge that traditionally would not have been recognized in a formal accounting of skills or technologies. Although the content of these categories is considered proprietary, Joseph Daniele, who oversaw the specific knowledge audit, describes them as "tricks, shortcuts, complex solutions" used in Xerox's product-development process.

The 28 knowledge categories were arrayed across product classes and competitors to form a matrix. Senior and middle-level technical managers were then interviewed and asked to assess qualitatively the state of Xerox's specific-knowledge holdings in each area--as well as to project Xerox holdings (and those of competitors) five and ten years into the future. Once analyzed, these data were made available to strategic planners as well as to Xerox's intellectual asset management function.

The specific-knowledge audit delivered clear benefits to both the strategic and intellectual-asset management functions. For example, it identified emerging knowledge gaps with respect to the paper-marking competency (e.g., an anticipated shortfall in some areas of specific knowledge due to a low "refresh" rate in the hiring of materials scientists). It also highlighted dangers in the voluntary withdrawal from some business segments (i.e., the implications of diminished capability in certain fast-moving areas of knowledge). The specific-knowledge audit helped Xerox to understand better its "human intellectual capital" by providing a language--beyond that of traditional skills terminology—for assessing what people in the firm know and how this knowledge relates to strategically important capabilities. Finally, the audit suggested ways in which specific knowledge could be more effectively leveraged. For example, it identified areas of Xerox operations where the transfer of highly tacit knowledge was critical and where, therefore, the co-location of knowledge "sources" and knowledge "recipients" should be considered.

Summary

Tacit knowledge has unique properties that confer power on management initiatives that leverage it effectively. First and perhaps most importantly, tacit knowledge tends to mirror the way work actually gets done within an organization. As such, it is powerful in the replication of effective practices, as reflected in the submissions redesign case. Second, tacit knowledge is created on the leading edge of the firm's learning; it reflects the things that people are learning but have been unable (or perhaps just too busy) to articulate and share with others. As such, leveraging tacit knowledge is key to adaptability in a rapidly changing environment--as it has been for the U.S. Army during a time of great change. Finally, as the Xerox case reflects, tacit knowledge thrives in the "white space" between official functions, processes, and product lines. As such, it can serve as a powerful lens through which to take stock of and manage a firm's intangible assets.

Working with Tacit Knowledge: General Approaches

Collect

For many firms, the point of entry to more deliberate knowledge management has been the attempt to collect and codify knowledge. With respect to tacit knowledge—particularly the “embodied” variety—this approach is sharply limited. Knowledge that people have trouble articulating tends to be time-consuming and expensive to codify (as witnessed by the knowledge-acquisition bottleneck in expert systems development). By contrast, tacit knowledge of the “embedded” variety can often be effectively leveraged through codification. Such knowledge may have remained tacit simply because no one thought it valuable enough to record and/or disseminate.

The terms “knowledge base” and “knowledge repository” have come into usage recently and it is worth taking a moment to consider what these terms denote, above and beyond a more familiar term such as “database.” When embedded knowledge is captured and stored, it tends to assume a greater variety of forms than do stored data or information. For example, data concerning the molecular weights and solubilities of the compounds in a series can be reliably represented as a table of scalar values--regardless of the particular compounds in question. By contrast, the record of a scientific team’s “lessons learned” may assume many forms, depending on whether these lessons concerned the selection of synthesis directions, the development of biological assays, or the management of project-team dynamics. In short, “knowledge bases” or “knowledge repositories” must handle content that is more heterogeneous and generally less structured than that stored in conventional databases. Because such content requires a greater degree of interpretation for its use, knowledge repositories tend also to be more closely linked to systems or functionality for communication and collaboration. The integration of knowledge repositories with messaging systems (as exemplified in groupware environments) reflects the fact that people often need to interact with one another directly in order to clarify, interpret, or validate the content that they find in a knowledge repository.

Knowledge repositories can help companies to derive value from embedded knowledge that might otherwise have “fallen through the cracks.” As such, they are becoming necessary adjuncts to more conventional forms of data and information storage. Knowledge repositories give people a place to record what they are learning (if they can be induced to do so) and they provide visibility and access to that knowledge across time and space. Knowledge repositories place relatively few constraints on the type of content they will accept and so are a natural choice to support a team, task force, or initiative whose information and knowledge requirements are still evolving. Add to these factors the advent of cheap, nearly ubiquitous client-server computing and it is not surprising that the development of knowledge repositories has become a focal point for knowledge-management efforts at many companies.

Yet knowledge repositories do not by themselves constitute solutions to the problem of identifying, capturing, and leveraging knowledge. Rather, they are best viewed as supporting elements in a solution that includes the familiar trio of technology, people, and processes. Repositories are most likely to deliver value when they are implemented as part of a systematic “lessons learned” process that includes defined activities, roles, and measures that together encourage knowledge capture and reuse. Perhaps the most important of these activities is the systematic debriefing of important events (e.g., project milestones, feedback from internal customers) and the capture of new knowledge that surfaces in a debriefing. Debriefings address the problem of getting people’s knowledge *into* the repository in the first place—the supply side of the internal knowledge market. Debriefings need not be extensive or complicated (as the experience of the U.S Army shows) but they must be connected to an effective system for bringing captured knowledge to bear on decisions. If “lessons learned” are seen as destined for someone’s filing cabinet alone, then both time and credibility will have been wasted.

Bringing the knowledge in a repository to bear on decision making requires attention to the demand side of the knowledge market as well. Users of the system (knowledge buyers) must come to expect that they will find material there that is comprehensible and useful to them in their work. If not, they will fail to incorporate the system into their working lives and the commitment to populating it with new knowledge will break down. To ensure quality and usefulness, some mechanism is typically required for warranting and qualifying captured knowledge—either “before the fact” (by vetting candidate knowledge before it is stored) or “after the fact” (by obtaining feedback from early users and using that feedback to validate, refine, or remove content). In either case, the demand side of the knowledge equation requires value-added reshaping of the documents and descriptors that populate the repository. Beyond considerations of quality or validity, documents often need to be restructured if they are to be of use in new settings and for new purposes.

All of which may amount, in a given organization, to making some very busy people even busier. Indeed, among the wide-ranging (and valuable) discussion of knowledge-sensitive measurements and incentives, the plain fact remains that many of the most valuable people in an organization cannot carry much more than they are already carrying. Quite naturally, they will resist any new system or initiative that they see as coming between them and the accomplishment of their primary business objectives. Because of this limit on the time and attention that knowledgeable people can devote to knowledge-management activities, firms within a number of industries (e.g., professional services, information technology, biopharmaceuticals) have begun to experiment with a new type of support role—the “knowledge integrator” or “knowledge coordinator.” Although the particulars vary, knowledge coordinators are generally responsible for collecting, reshaping, and disseminating the information and knowledge that others in the firm produce in the course of their work. Knowledge coordinators also function as intermediaries, putting knowledgeable people in touch

with those who need to draw upon their knowledge—usually within a limited domain of practice (e.g., an industry vertical in a consulting firm). With talented professionals busier (and more expensive) than ever, a business case appears to be emerging for using skilled knowledge coordinators to leverage the professional and specialist knowledge of the firm. Although we are only beginning to understand it¹³, the concept of a dedicated knowledge coordinator seems to represent a formalization of the value-adding activities of informal knowledge brokers as well as an acknowledgment of the limitations of purely informal mechanisms to meet the knowledge demands of the modern enterprise.

As stated earlier, attempts to codify tacit knowledge of the “embodied” variety will tend to yield disappointing results. At times, however, an intensive assault on embodied knowledge is required. For example, when the concentration of mission-critical knowledge in one or a few experts exposes the firm to unacceptable levels of risk, then knowledge-engineering approaches may be in order. Focused interviews, think-aloud protocols, and sorting/mapping exercises can be used to elicit and codify previously tacit knowledge in order to populate a rule-based expert system. Libraries of problems/cases can be generated, analyzed, and indexed for use in a case-based reasoning system. Classifier systems (such as those based on neural-network simulations) can be used to induce and recognize the complex patterns that underlie expert decisions. In general, investment in these intensive approaches to knowledge codification will produce the greatest returns when the decision-support tools that they yield can be widely deployed and when the knowledge that they encode is relatively slow to “turn over.”¹⁴

Connect

At IBM Global Services, we often tell consulting clients that if they have a dollar to spend on knowledge management, they will get the most from that dollar if they spend it connecting people to one another rather than attempting to codify what those people know. The primary argument, as we have said, is that much of the valuable knowledge in an organization is embodied, context-dependent, subject to a short shelf-life and otherwise “sticky.” This sticky, tacit knowledge loses something in the translation/codification. Often it is preferable to bring the problem to the knowledgeable person rather than trying to bring his or her (disembodied) knowledge to the problem. A second, more philosophical argument is that all knowledge is socially constructed and, by its nature, not fully separable from the social interactions through which it is developed. According to this argument, knowledge is transferred through interactions that are rich and grounded in shared values and warrants; thus knowledge management should focus primarily on increasing the frequency and quality of those interactions.

Whatever the rationale, an important approach to working with tacit knowledge is to increase the connectedness of the organization, defined here as the ability of knowledge buyers and knowledge sellers to find one another and interact to some productive end. Because global organizations are huge and dispersed—and because personal expertise is often opaque—the problem of “knowing what we know” must

be attacked on several fronts. One line of attack is through the provision of “people finding” technologies ranging from corporate directories to sophisticated, adaptive systems for expertise profiling and referral. A related, often prior, line of attack is through a knowledge-mapping exercise in which relevant areas of expertise are identified, a taxonomy of descriptors is developed, and a process is put in place to update and maintain the set of expert profiles. A third line of attack, to which we have already alluded, is the use of human intermediaries or “brokers” who put knowledge buyers and knowledge sellers in touch with one another.

The knowledge and expertise of individuals within the firm can be identified, verified, and stored in browseable/searchable form--yet what the firm knows may still be incompletely represented. This is because much of the knowledge of the firm is embodied not in individuals but in the informal social networks and communities in which individuals participate. Some of the most powerful approaches to increasing connectedness (and thus leveraging tacit knowledge more effectively) are those focused on identifying and cultivating these networks, which seldom correspond to the portrait of the official organization. A great deal of research attention has been paid to communities of practice and the practical fruits of this work are beginning to ripen in the form of methodologies for uncovering and strengthening communities. Variants of social network analysis can be used to map the flow of information and knowledge within a firm and, in so doing, uncover communities-of-practice. Once uncovered and understood, communities can be strengthened in a variety of ways, such as improving their visibility and recognition, providing support for community-building events and activities, and providing tools (websites, workrooms) to allow community members to interact easily and asynchronously.

Finally, however, a limit may be reached in how effectively knowledge can be brought to bear simply by putting people in touch with one another. Even when they are willing members of a community, experts have personal business objectives against which they must deliver and, as such, must limit the degree of time and mental energy they can give to someone else’s problem. Further, experts embedded in social networks still leave the firm vulnerable to their loss. In either event, more intensive (and intrusive) approaches to connection can be found in traditional knowledge-transfer paradigms such as mentoring and apprenticeship, accelerated development through job assignments, and outside recruitment/acquisition.

Characterize

When compared to the attention paid to collecting knowledge and connecting knowers, efforts to characterize tacit knowledge for the purpose of managing it more effectively have been somewhat neglected. This is unfortunate, for many firms launch into the “how” of working with tacit knowledge without fully understanding the “what” or even the “why” of working with tacit knowledge. In our consulting practice, we have found that the characterization of knowledge assets is critical to deriving full value from knowledge-

management programs. A variety of approaches to characterization can most easily be summarized in terms of the familiar distinction between “bottom-up” and “top down” reasoning.

Bottom-up approaches begin with a description of how work gets done “on the ground”—the activities and decisions that make up a business process--and seek to arrive at broad classes or domains of knowledge that underlie and support the work. Critical to the bottom-up approach is a relentless focus on how things *actually* happen in the business--the informal practices and “work arounds,” the reliance on judgment calls and “sneaker networks”, and the political and cultural realities that remain unspoken. Through an iterative process of storytelling, decision mapping, and knowledge classification, models are constructed which depict work practices in terms of the knowledge assets that underlie them. These models can then be used to support management decision making across a wide range of issues; from deriving information and knowledge requirements for IT planning to evaluating and improving the way in which global projects are staffed.

Top-down approaches, by contrast, begin with the identification of broad knowledge domains—those determined to be of highest-value to the business—and seek to decompose those domains to the point where specific interventions (i.e., pilot projects) can be identified, scoped, and prioritized. Top-down approaches begin with strategic business issues rather than business activities/processes and, as a consequence, they tend to be undertaken with more senior management groups, typically in a workshop setting. Beyond simply identifying those areas of the business where more effective knowledge management can deliver value, top-down approaches to knowledge characterization are useful in building awareness of knowledge as a leverageable asset and in defining concrete opportunities for action.

Conclusion

In this article, I have explored some of the conceptual and practical issues that arise when companies seek to capture more business value from the hidden or tacit knowledge of their people. Beginning from a simple, working definition of knowledge, I have distinguished two forms of tacit knowledge (embodied and embedded) and presented a general case for working with both forms. To show how real companies have benefited from working deliberately with tacit knowledge, I have presented three brief cases and have noted their success factors. To show the range of possible solutions, I have reviewed some general approaches to working with tacit knowledge (connect, collect, and characterize) with the qualification that most successful programs will incorporate elements of all three approaches.

Running throughout this article is the proposition that the knowledge-based organization needs to “know what it knows” and cannot limit this knowing to the formal, codified knowledge that is the residue of past experience and analysis. Instead, companies that seek to win in the era of knowledge-based business will have to grapple with the hidden or tacit dimension of knowledge--the dynamic, leading edge of what they

know--in order to learn, innovate, and compete more effectively. Clearly, however, companies have always worked with tacit knowledge--socializing embodied knowledge through apprenticeships, codifying embedded knowledge in "tip sheets" and practitioner's handbooks, or trying to achieve the right mixture of skills and experience in departments and on project teams. If tacit knowledge has always been with us, do we need to make of it yet another management subdiscipline?

The history of management thinking shows that solutions to seemingly intractable problems often have awaited the development of an organizing concept where none existed before. This process, akin to what philosophers call "reification," amounts to taking a set of previously unconnected phenomena and treating them as instances of a single concept. In business settings, reification can set the stage for quantum improvements because it allows companies to begin to work deliberately (rather than incidentally) on problems that have a common root. Before the notion of a cross-cutting business process was introduced, for example, companies could work to improve the efficiency and responsiveness of their operations, but they could not work deliberately to redesign their processes. Today, companies need to work deliberately to improve their management of knowledge and they cannot be content with the informal and disconnected approaches of an earlier, industrial era. On this view, the value of tacit knowledge as a management concept lies in its potential to draw together and illuminate a set of business problems that --if not new--are newly important.

¹ Brown, J. S. & Duguid, P. "The Knowledge Continuum," unpublished manuscript. Nonaka, I. & Takeuchi, H. *The Knowledge-Creating Company: How Japanese Companies Create the Dynamics of Innovation*. (New York, NY: Oxford University Press, 1995). Leonard, D. & Sensiper, S. "The Role of Tacit Knowledge in Group Innovation," *California Management Review*, Vol. 40/3 (Spring 1998): 112-132. Sternberg, R. J. & Horvath, J. A. *Tacit Knowledge in Professional Practice*. (Hillsdale, NJ: Lawrence Erlbaum Associates, in press).

² Polanyi, M. *The Tacit Dimension*. (New York, NY: Doubleday, 1966). Reber, A. *Implicit Learning and Tacit Knowledge*. (Boston, MA: Cambridge University Press, 199_).

³ Nelson, R. & Winter, S. *An Evolutionary Theory of Economic Change*. (Cambridge, MA: Harvard University Press, 1982).

⁴ Leonard-Barton, D. *Wellsprings of Knowledge: Building and Sustaining the Sources of Innovation*. (Boston, MA: Harvard Business School Press, 1995).

⁵ Pisano, G. *The Development Factory*. (Boston, MA: Harvard Business School Press, 1997).

⁶ Nonaka, I. & Takeuchi, H. *The Knowledge-Creating Company: How Japanese Companies Create the Dynamics of Innovation*. (New York, NY: Oxford University Press, 1995).

⁷ O'Dell C. & Grayson, C.J. *If Only We Knew What We Know: The Transfer of Internal Knowledge and Best Practice*. (New York, NY: The Free Press, 1998).

⁸ Szulanski, G. "Exploring Internal Stickiness: Impediments to the Transfer of Best Practice within the Firm," *Strategic Management Journal*, 17 (Winter Special Issue, 1996): 27-43.

⁹ Winter, S. "Knowledge and Competence as Strategic Assets," in D. Teece, ed. *The Competitive Challenge* (New York, NY: HarperCollins, 1987): 159-184.

¹⁰ Kogut, B. & Zander, U. "Knowledge of the Firm: Combinative Capabilities and the Replication of Technology," *Organization Science*, 3/3 (August, 1992): 34-40.

¹¹ Leonard-Barton, D. *Wellsprings of Knowledge: Building and Sustaining the Sources of Innovation*. (Boston, MA: Harvard Business School Press, 1995).

¹² Hamel, G. & Prahalad, C. *Competing for the Future*. (Boston, MA: Harvard Business School Press, 1994).

¹³ At the IBM Institute for Knowledge Management, we are currently undertaking a study of emerging best practices in this area.

¹⁴ Davenport, T. & Prusak, L. *Working Knowledge: How Organizations Manage What They Know*. (Boston, MA: Harvard Business School Press, 1998).