

Doing Knowledge Management

By

Joseph M. Firestone
Executive Information Systems, Inc., Alexandria, VA

and

Mark W. McElroy
Center for Sustainable Innovation, West Windsor, VT

Introduction

Has Knowledge Management (KM) been done? Of course, KM has been done. It is a natural function in human organizations, and it is being done all of the time in an informal distributed way by everyone undertaking activity in order to enhance knowledge production and integration tasks. But whether formal interventions claiming the label "KM" are bona fide instances of KM practice is another matter entirely. To answer that question, we need to have clear, non-contradictory ideas about the nature of knowledge, knowledge processing, and Knowledge Management. And to have those, we need *to get beyond the notion that we can do KM by just doing anything that may have a positive impact on worker effectiveness while calling that thing "KM."*

Instead we need to recognize that *the immediate purpose* of KM is *not* to improve either worker effectiveness (though it may well do that) or an organization's bottom line. Its purpose is *to enhance knowledge processing* (Firestone and McElroy, 2003, ch. 3) in the expectation that such enhancements will produce better quality solutions (knowledge), which, in turn, *may, ceteris paribus, when used*, improve worker effectiveness and the bottom line. And when we undertake KM projects, we must evaluate the contributions of our interventions to the quality of knowledge processing and knowledge outcomes. That calls for tough, precise thinking about knowledge processing, knowledge, and the impact on these that our interventions are likely to have.

The question we are asking here is whether KM practitioners are, in fact, providing this tough, precise thinking as a basis for KM practice, or whether, instead, they are "practicing KM" by helping fields or techniques such as Information Technology, Content Management, Customer Relationship Management (CRM), Social Network Analysis, Storytelling, Communities of Practice, and "Knowledge" Cafés to "colonize" it? Is such conceptual drift in KM so widespread that one can conclude that, generally speaking, at least, *KM as a formal, intentional endeavor* has, indeed, not yet been done?

In this paper we will begin by providing an account of our view of KM, knowledge processing, information, knowledge, and Knowledge Management, and then continue by considering the above questions and by analyzing the Partners HealthCare case, a case where KM has most emphatically been done, and done successfully. We will then end by drawing out the implications of the Partners HealthCare case for KM Strategy and KM Programs.

The Nature of KM as a Type of Activity or a Set of Processes

In an earlier "Viewpoint" in TLO (Firestone and McElroy, 2004) we presented a three-tier framework (see Figure 1) of business processes and outcomes (Also see McElroy, 2003, Firestone, 2003, and Firestone and McElroy, 2003, 2003a), distinguishing *operational business processes*, *knowledge processes*, and *processes for managing knowledge processes*. *Operational processes* are those that use knowledge but, apart from routinely produced knowledge about specific events and conditions, don't produce or integrate it. Examples of outcomes are Sales Revenue, Market Share, Customer Retention and Environmental Compliance.

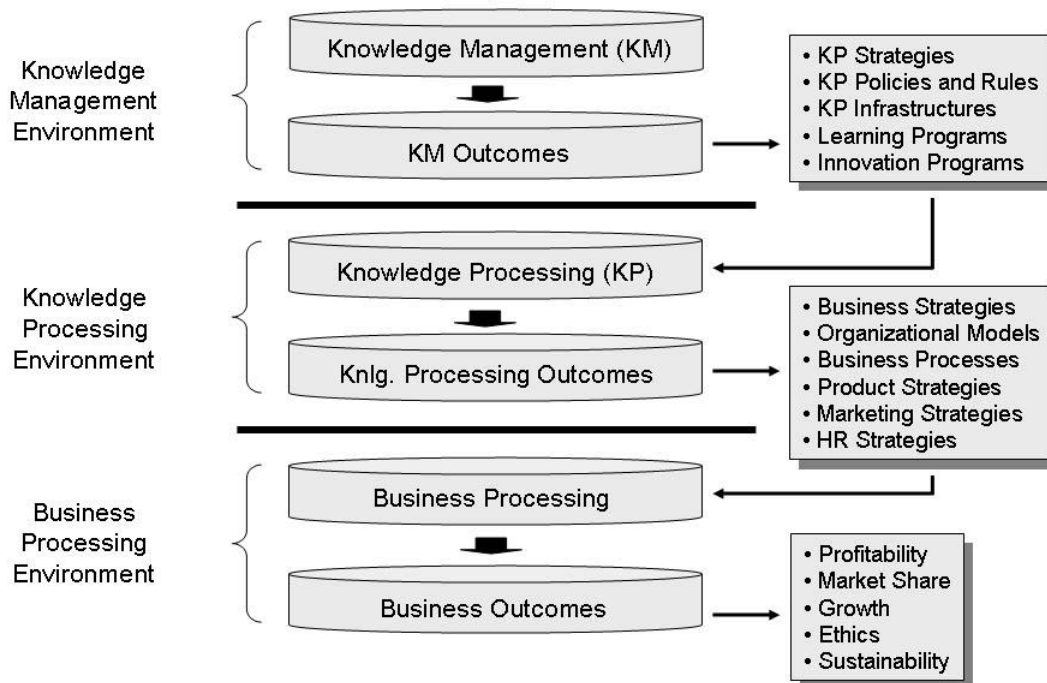


Figure 1. The Three-Tier Framework

There are two *knowledge processes*: knowledge production, the process an organization executes that produces new general knowledge and other knowledge whose creation is non-routine; and knowledge integration, the

process that presents this new knowledge to individuals and groups comprising the organization. Examples of outcomes are new organizational strategies communicated throughout an enterprise using e-mail, and new health insurance policies communicated through a new release of the organization's personnel manual.

Knowledge Management is the set of processes that seeks to change the organization's present pattern of knowledge processing to enhance both it and its outcomes. A discrete Knowledge Management activity is one that has the same goal as above or that is meant to contribute to that set of processes. The discipline of KM is the study of such processes and their impact on knowledge and operational processing and outcomes. The foregoing implies that KM doesn't directly manage, create or integrate most knowledge outcomes in organizations, but only impacts knowledge processes (performed by operational process agents), which, in turn, impact knowledge outcomes. For example, if a Knowledge Manager changes the rules affecting knowledge production, then the quality of knowledge claims may improve. Or if a KM intervention supplies a new search technology, based on semantic analysis of knowledge bases, then that may result in improvement in the quality of business forecasting models.

The Context of KM: CASs, DECs, and Learning

What is the conceptual context of this three-tier conceptualization of KM? It is in the integration of the theories of Complex Adaptive Systems (CASs) (Holland, 1995, Gell-Mann, 1994, Kauffman, 1995, Juarrero, 1999, Hall, 2005) and Organizational Learning (OL) (Argyris and Schon, 1974, Argyris, 1993, and Senge, 1990). The three types of processes distinguished in the three-tier framework occur within complex adaptive organizational systems that are characterized by distributed continuous learning and problem solving, self-organizing, and emergent phenomena produced by dynamic processes of interacting autonomous agents that are non-deterministic in character (Holland, 1998). Emergent phenomena at the group and global system levels in organizations exhibit "downward causation" on individual decision makers in such systems (Campbell, 1974, Bickhard, 2000). These phenomena include social, geo-physical, economic, and cultural conditions, and also social network effects presented to individuals in the form of transactions directed at them by other decision makers who collectively constitute the emergent network pattern (see Figure 2) of the organizational CAS (Firestone and McElroy, 2003, chs. 2 and 4).

When we look more closely at individual CAS agents and their decisions, we connect to matters that have received a great deal of attention in the field of organizational learning. Decisions are part of a sequence of cognitive operations that have been described in the literature in slightly varying terms, using many names (e.g., the organizational leaning cycle, Ackoff, 1970, the experiential learning cycle, Kolb and Fry, 1975, Kolb, 1984, the adaptive loop, Haeckel, 1999,

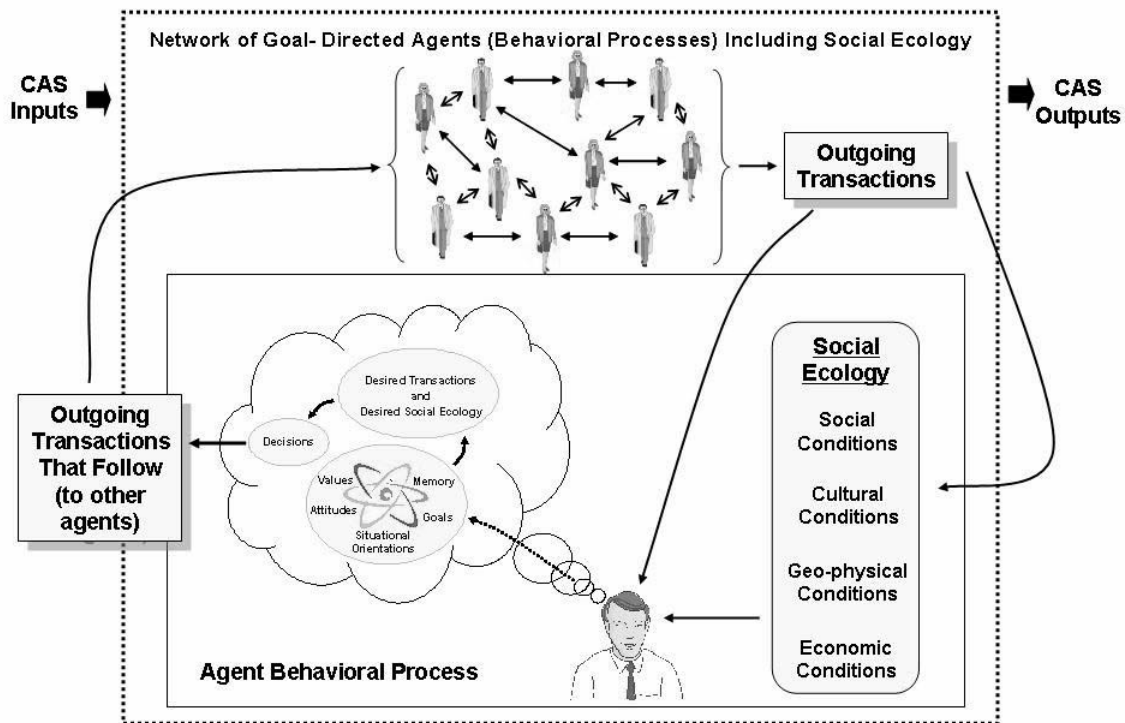


Figure 2. The Organizational CAS

and others). We call it the Decision Execution Cycle (DEC), which includes Planning, Acting, Monitoring, and Evaluating behaviors (Firestone, 2000). Decisions are produced by planning and are embodied in acting. Decisions produce actions. And actions - activities - are the stuff that social processes, social networks, and (complex adaptive) organizational systems are made of. Figure 3 illustrates the phases of Decision Execution Cycles.

DECs use previously existing individual-level knowledge to arrive at decisions and actions. Personal knowledge is always the immediate precursor to action. DECs also generate new knowledge about specific conditions and situations by using preexisting knowledge in a routine way to monitor, evaluate, plan and decide. This is the Single-Loop Learning (SLL) of Argyris and Schön (1974). In addition, DECs play a key role in initiating and performing Double-Loop Learning (DLL) (Argyris and Schön, 1974) – learning of new knowledge (in the form of general predispositions and rules, and specific knowledge) that requires problem solving and is not just a matter of perception or direct apprehension or comprehension.

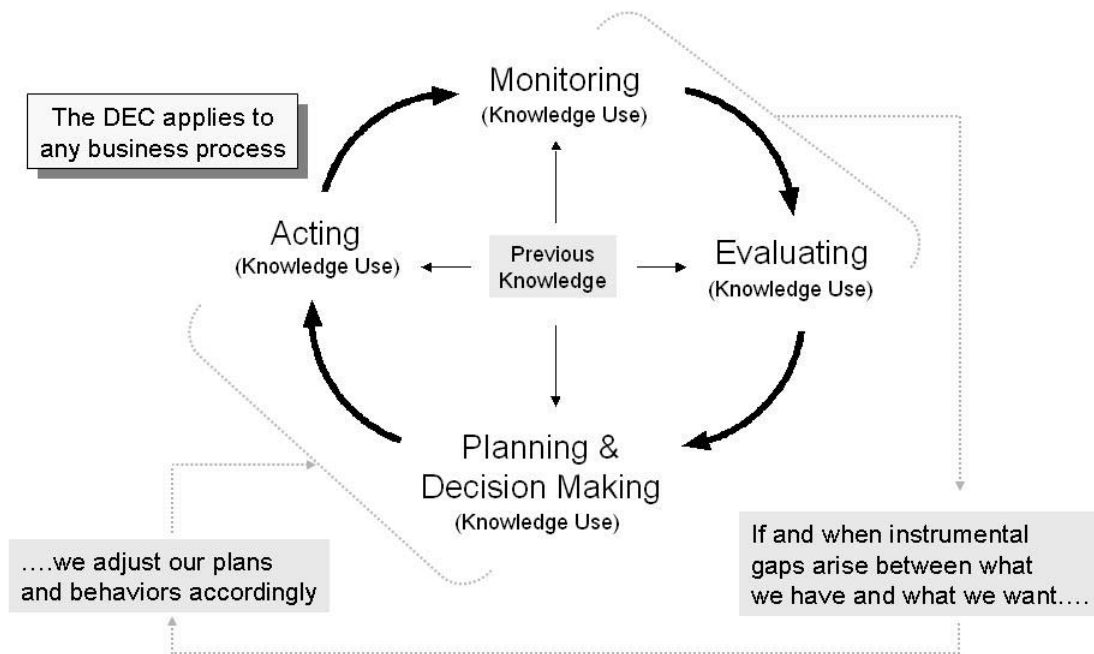


Figure 3. The Decision Execution Cycle

Elsewhere, we (Firestone and McElroy, 2003, 2003b and Firestone, 2003, 2003b) have described how routine DEC's give rise to DLL. In brief, DEC decisions and actions are accompanied by expectations. During monitoring and evaluating, the individual determines the degree to which results match the expectations accompanying decisions, and when mismatches occur, the seriousness of the mismatch from both the factual and evaluative perspectives (see Figure 4). When the mismatch is great enough from the viewpoint of the individual, and when the individual decides that previous knowledge won't work to reduce the mismatch, the individual recognizes that a gap exists between what the individual knows and what she or he needs to know in order to pursue the goal(s) or objective(s) of the associated DEC's. This knowledge, or epistemic, gap is what we mean by a "problem," and recognition of it is what we mean by "problem recognition."

When a DEC results in problem recognition, the individual can either abandon or suspend pursuing the goal or objective motivating associated DEC's or alternatively, the individual can engage in problem solving or DLL, a process composed of multiple learning-related DEC's motivated by a learning incentive. Following Popper (1999), we view DLL most generally as an emergent (i.e., non-deterministic) three-stage knowledge process comprised of problem formulation, developing alternative solutions, and error elimination, the stage in which we

select among alternatives by eliminating the ones we think are false. Among the results of error elimination is *knowledge*, which we'll discuss briefly below. Here we call attention to the need, once new knowledge is produced, for further knowledge processing to integrate it into the DEC and business process environment that originated it, and into the organizational memory that will make it available for re-use later.

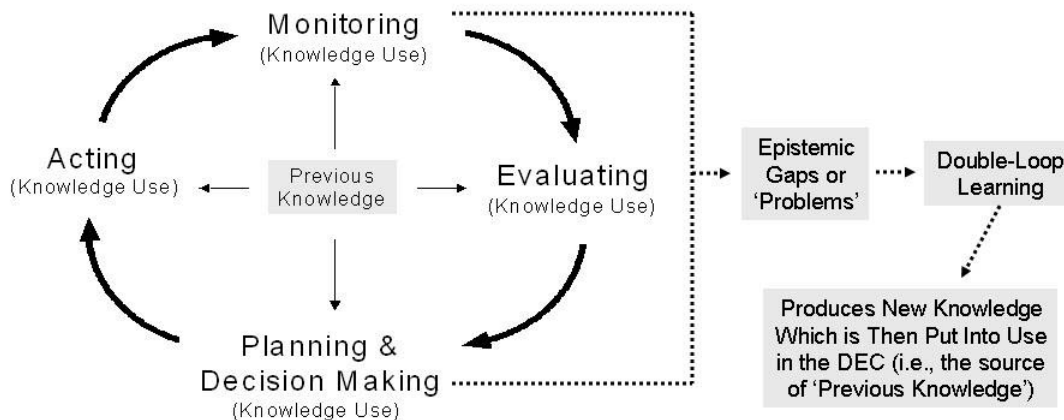


Figure 4. The Decision Execution Cycle and Problem Recognition

The Knowledge Life Cycle, the Business Processing Environment, and the DEC

So far, our account of DLL/problem solving as involving sequences of DEC's has focused on the individual level of analysis. But DEC's may also form patterns of interpersonal collaboration, cooperation, and conflict, and these patterns may also integrate into knowledge processes. When they do, we can differentiate between problem formulation, developing alternative solutions, and error elimination, on the one hand, and problem *claim* formulation, knowledge *claim* formulation, and knowledge *claim* evaluation in order to distinguish the individual level of knowledge processing from the interpersonal and collective levels, respectively. We also distinguish information acquisition and individual and group learning, as additional knowledge sub-processes preceding knowledge claim formulation. Information acquisition includes activities of finding and retrieving

knowledge claims produced in external systems. Individual and group learning is a category identifying levels of knowledge processing nested within the knowledge production process being analyzed. Individual and group learning produces knowledge from the viewpoint of nested knowledge processes, and knowledge claims from the viewpoint of knowledge claim formulation at higher levels of analysis.

When we view knowledge processing at levels of analysis higher than the individual level, we identify the pattern including problem claim formulation, information acquisition, individual and group learning, knowledge claim formulation, and knowledge claim evaluation as the knowledge production process resulting in both new tested and surviving beliefs and knowledge claims. Once new knowledge is produced at the collective level, it must be integrated into organizational memory, key DEC's and business processes. This process of knowledge integration is made up of four more sub-processes, all of which may use interpersonal, electronic, or both types of methods in execution. They are: knowledge and information broadcasting, searching/retrieving, knowledge sharing (peer-to-peer presentation of previously produced knowledge), and teaching (hierarchical presentation of previously produced knowledge).

Knowledge integration is about system-level knowledge claims being communicated from one part of the Distributed Organizational Knowledge Base (DOKB), the configuration of previously produced knowledge claims, beliefs and belief predispositions in the organization (Firestone and McElroy, 2003) (see Figure 5), to another. Knowledge claims are stored in media and information systems. Beliefs and belief predispositions are stored in minds. Through the DOKB, both knowledge claims and belief phenomena are accessible in varying degrees to individual decision makers in DEC's, within both the Business Processing Environment, and the knowledge and KM processing environments. That is, *the DOKB is the knowledge and information foundation for all of the organization's DEC's and processing environments*. When knowledge claims are evaluated, results of evaluation in the form of changes in beliefs and new knowledge claims, including those we call "meta-claims" which provide the "track record" of criticism, testing, and evaluation of knowledge claims produced during knowledge claim formulation, are stored in the DOKB. *Knowledge claims, as well as meta-claims, are then integrated and reintegrated into the DOKB as they are broadcasted, retrieved, shared and taught again and again.*

A visual of knowledge processing and its relationship to operational business processing is given in Figure 6, the Knowledge Life Cycle (McElroy, 1999, 2000, 2003, Firestone, 2000, 2003a, Firestone and McElroy, 2003, 2003a, 2003b, Cavaleri and Reed, 2000, 2001). Actually, the KLC extends from problem claim formulation to the integration of knowledge and information in the DOKB. Knowledge claim evaluation (KCE) occupies a central place in the visual and in knowledge production. It is KCE that produces surviving, falsified, and undecided knowledge claims, and also meta-claims, for storage in the DOKB. Of course, the extent to which this "track record" is stored or lost depends on the specifics of each organization. The bottom of the figure illustrates the workings of the

business processing environment, including its role in using knowledge for business processes and in recognizing problems that arise through mismatches of results and expectations, which, in turn, initiate DLL/knowledge production activity.

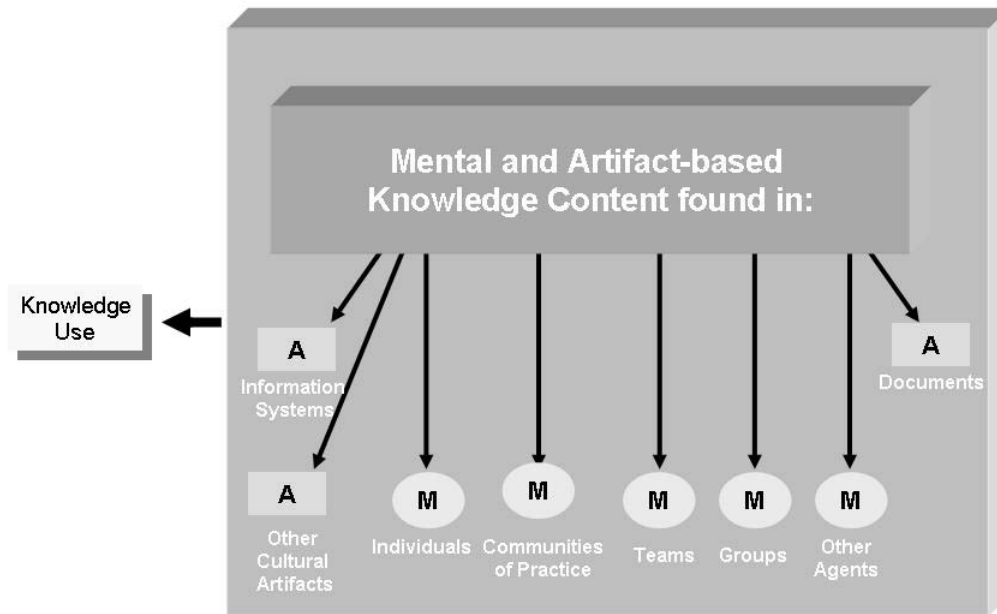


Figure 5. The Distributed Organizational Knowledge Base (DOKB)

The clouds in the figure illustrate the ubiquity of DOKB content in the various processes. We have also used arrows from the primary DOKB cloud to illustrate its influence on all processes, but are limited to showing its universal influence in two dimensions, while at the same time showing the breakdown of primary knowledge processes into sub-processes and other details in the figure.

Since Figure 6 focuses on a process view, it glosses over the lower DEC level of analysis. Figure 7 makes it clear that the match/mismatch process occurs in DEC's and not simply at the higher level of business processes. This point is very important for our later analysis of the Partners HealthCare case.

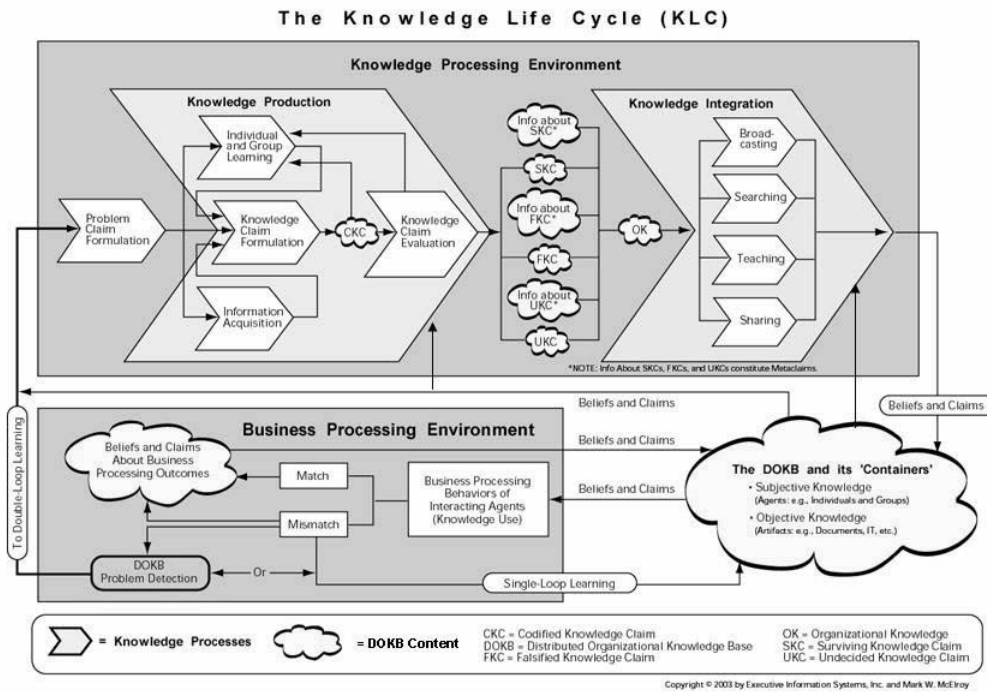


Figure 6. The Knowledge Life Cycle and the Business Processing Environment

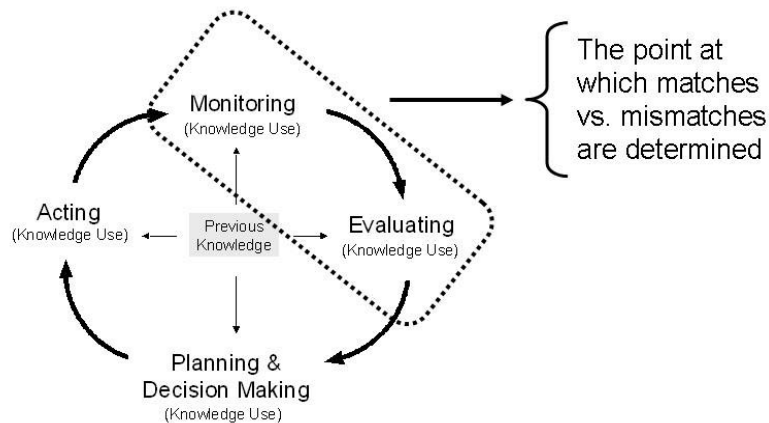


Figure 7. Matches and Mismatches and the DEC

Information and Knowledge

Information is a non-random structure within a system, *indicating future interactive potentialities*, either originating along with it, or acquired or developed by it in the course of its interacting with and responding to its environment and the problems generated by that interaction (Bickhard, 1999). Note that this definition does not require correspondence between information and the environment. Nor does it assert that information is encoded in some simple cause-and-effect fashion, but leaves room for emergent information in the context of interaction with the environment.

The most important aspect of information, in our view, however, is not whether it is complex or simple, or produced quickly or slowly, or gained or lost over time, or whether there is a great or a small amount of it. All of these are undoubtedly important, but the most important aspect of information is whether its influence on behavior enhances the ability of the system using it to adapt. And this ability to adapt, in turn, is most likely to be enhanced if the information itself actually corresponds to the reality of the system's environment. Evolution provides such correspondence by selecting for those life forms that fit the environmental constraints in which they live. Errors in genetic information are eliminated over time by the environment, along with the organisms that contain them (Popper, 1987). Learning provides such correspondence on a much shorter time scale by providing us with an opportunity to eliminate our errors in information and to create new information that survives our evaluative efforts and our experience.

Since the most important aspect of information is correspondence with reality, the most important measures of information networks are those that evaluate this correspondence. Thus, the most important measures we can develop describing knowledge claim (information) networks are measures that help us to evaluate knowledge claims, and that brings us to "knowledge." One of the moments of truth in any consideration of KM is when it is time to say what one means by "knowledge." We favor a "unified theory" that specifies a viewpoint about the general phenomenon, but which also distinguishes different types of knowledge.

Knowledge is a tested, evaluated and surviving structure of information (e.g., DNA instructions, synaptic structures, beliefs, or claims) that may help the living system that developed it to adapt. This is our general viewpoint. It is consistent with our definition of information. And it is consistent with CAS theory and the view that knowledge is something produced by CASs in order to help them adapt to environmental challenges.

There are three types of knowledge:

- Tested, evaluated, and surviving structures of information in physical systems that may allow them to adapt to their environment (e.g., *genetic* and *synaptic* knowledge)

- Tested, evaluated, and surviving beliefs (in minds) about the world (subjective, or non-sharable, *mental* knowledge)
- Tested, evaluated, and surviving, sharable (objective), linguistic formulations about the world (i.e., claims and meta-claims that are speech- or artifact-based or *cultural* knowledge)

The ontology reflected in the above definition is from Popper (1972, 1978, 1994, 1999, Popper and Eccles, 1977), but we have not used his terminology here. Figure 8 illustrates the three types of knowledge and depicts their abstract relationships.

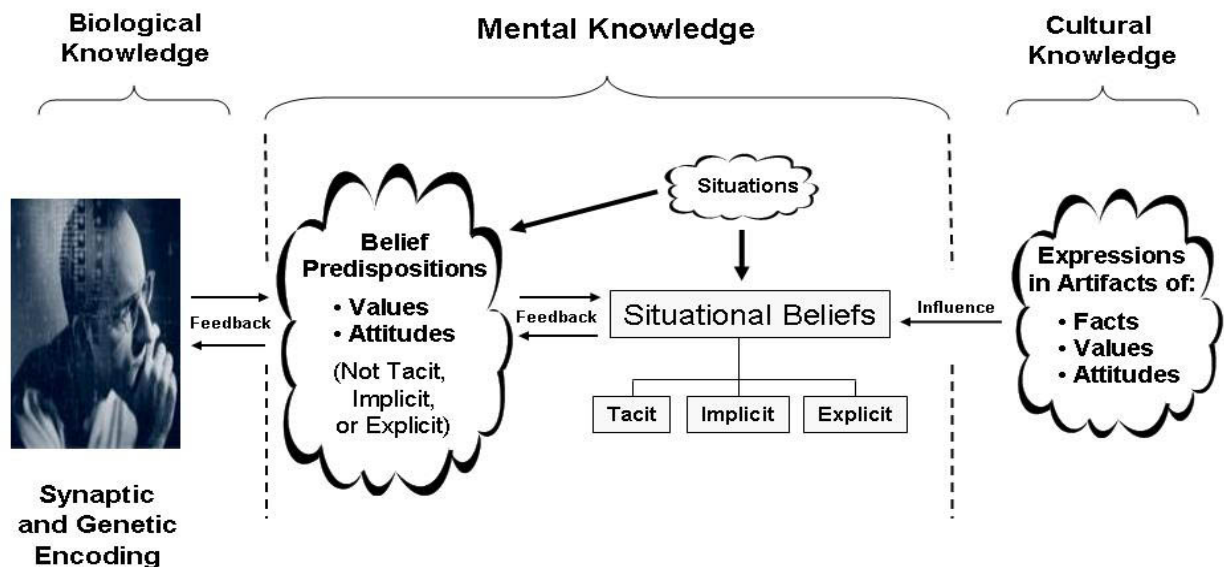


Figure 8. The Three Types of Knowledge

Doing KM?

At the beginning of this paper, we raised two questions related to the theme of this special issue. Are KM practitioners "doing KM," or are they "practicing" KM by helping fields or techniques such as Information Technology, Content Management, CRM, Data Warehousing, Social Network Analysis, Storytelling, Communities of Practice, Data Mining, Quality Management, Human Resources, and "Knowledge" Cafés to "colonize" it? Is conceptual drift in KM so widespread that one can conclude that, generally speaking, at least, *KM as a formal, intentional endeavor* has, indeed, not yet been done? The detailed answers to

these questions depend on one's conceptual orientation to KM. Now that we've laid that orientation out, we can offer an analysis that will provide some answers.

First, we do think that KM as a formal intentional endeavor has been done, and later we will discuss a case study that will illustrate this in detail. Having said that, however, we also believe that far too many "KM" efforts are not KM at all, but represent activities in fields peripheral to KM that "colonize" it by using KM terminology to mischaracterize non-KM interventions as instances of KM with the intention of benefiting from its cachét.

Second, in fact, such colonization of KM is not new. KM has been subject to it from the beginnings of the discipline, when it was frequently characterized as being about "delivering the right information to the right people at the right time," through use of the right IT tool. Thus, KM was viewed as an activity that encompassed deploying the right IT tool in the enterprise and, often, using it to "manage knowledge" as characterized above.

In that spirit, Data Warehousing, Data Mining, Business Intelligence (BI) and Online Analytical Processing (OLAP), Business Performance Measurement (BPM), CRM, Enterprise Resource Planning (ERP), Collaboration Management, Groupware, Search and Retrieval applications, Content Management (CM), Semantic Network/Text Mining applications, Document Management, Image Management, e-Conference applications, e-Learning applications, Expertise Locators (Yellow Pages), Best Practices Database applications, and Enterprise Information Portals (EIPs), have all been characterized as KM tools, and projects involving the deployment and use of one or another of these tools have been characterized and reported as KM projects. EIPs, in fact, were characterized as KM's "killer app," and scores of "KM cases" involving EIP projects were described and analyzed in the KM and portal literature (Firestone, 2003a).

At present, it is commonplace for portal vendors to characterize their search and retrieval capabilities as KM capabilities, as if in using them one was automatically managing "knowledge" and also finding it. Of all of the above IT applications, the most widely deployed in the 1st generation of KM was the Best Practice Database application.

Third, in our view, the association of the idea of "KM intervention" with any of the above tools is frequently an instance of "conceptual drift," mistaking KM for other forms of activity. Such drift is harmful to KM because, ultimately, it confuses the record of KM performance and therefore prevents an evaluation of KM based on that performance. Thus, because of conceptual drift it is possible to say that KM projects have failed 85% of the time, when, in fact, neither KM interventions, nor an evaluation of them in any quantity, may actually have occurred.

But how can one tell in any individual case of a project or program, that it is, in fact, a KM intervention, rather than an intervention of another type? The answer is that (a) one must evaluate an intervention with one's ideas about KM, knowledge processing, knowledge, information, other business processes and

outcomes, and the differences among them in mind, and (b) one must evaluate an intervention as if its classification as a KM intervention is conjectural and must be evaluated against alternative conjectural classifications. Of course, the quality of one's evaluation will be dependent upon the quality of one's KM framework, and also on the extent to which one has considered other alternative classifications for the intervention.

Fourth, the situation with respect to analytical or social interventions is quite analogous to that of IT tools. As KM developed, Communities of Practice (Wenger, 1998) became a popular, even dominant, "KM" intervention. Soon it was supplemented with storytelling (Denning, 2001) interventions encouraging knowledge workers to use stories to both "sell" KM internally, share knowledge, and facilitate collaboration. More recently, Social Network Analysis (SNA) (Cross and Parker, 2004) is being used to discover the structure of relationships in existing communities, as well as the existence of clusters of social relationships that can form the nuclei of new communities not yet self-organized.

Another technique that has been popular is the Knowledge Café (Isaacs, 1999), a technique in which participants circulate among multiple small interactive groups carrying on a discussion of a selected topic and sharing their knowledge over the course of a day. Additional techniques include "Knowledge" Auditing and Mapping, Value Network Analysis (Allee, 2003), Group Decision Making Processes, Influence Network Analysis, various Quality Management techniques, and, of course, Cultural Analysis.

Our view of this list of techniques is analogous to our view about IT tools. Specifically, projects or programs that use them are not automatically, or even by presumption, KM projects or programs. Whether they are, or not, depends on how the specific intervention is related to KM, knowledge processing, information, knowledge, and so on, and also on how it is related to other management and knowledge processing activities.

Fifth, "KM" interventions will involve either IT tools or social techniques or some mix of them. Whether such an intervention is a bona fide KM intervention depends on whether it is a policy, program, or project targeted at enhancing knowledge processing and through knowledge processing, knowledge outcomes, and ultimately business decisions and processes. In other words it depends on whether, and on the extent to which, the intervention fits the pattern expressed in the three-tier framework (see Figure 1), and is targeted at the KLC (see Figure 6) as compared with the extent to which it fits the pattern characteristic of other forms of management activity.

These considerations suggest that we apply the following criteria in deciding the question of whether an intervention is a KM intervention or something else:

1. Is the intervention aimed at having an impact on problem recognition in DEC's and business processes, on the KLC, or on some aspect of KM itself?

2. If the intervention is aimed at some aspect of knowledge integration in the KLC, or the DOKB itself, does it incorporate a way of telling the difference between knowledge and information so that its impact is aimed at knowledge integration and not just at information integration?
3. If the intervention is aimed at enhancing information acquisition relevant to a problem, does it incorporate a way of telling whether external information is or is not relevant to the problem?
4. If the intervention is aimed at enhancing knowledge claim formulation, does it incorporate tools or techniques for enabling creation of alternative knowledge claims?
5. If the intervention is aimed at knowledge claim evaluation, does it incorporate tools or techniques that will enable testing and evaluation of knowledge claims?
6. If the intervention is aimed at individual and group learning, does it meet any of the foregoing criteria about problem recognition, knowledge integration, or any of the knowledge production sub-processes?
7. If the intervention is aimed at enhancing KM itself, does it incorporate tools or techniques that facilitate any of the following: (a) any aspect of producing or integrating KM-level knowledge; (b) problem recognition in KM-level DEC's or business processes; (c) KM-level leadership; (d) building external relations with others in KM; (e) KM-level symbolic representation; (f) changing knowledge processing rules; (g) crisis handling in KM; (h) negotiating for resources; (e) allocating KM resources?

Sixth, while we don't have the space to apply these criteria to all of the techniques and tools listed above, we will apply them to a few of the most visible "KM" tools, techniques, and interventions. In the early days of KM, the most popular intervention was the development of Best Practice Database applications. The simple idea behind this type of solution is that the quality of decisions will improve if "Best Practices" are captured, made available to knowledge workers, and reused by them.

Are "Best Practices" interventions instances of "KM"? According to the criteria we've specified above they are not, because while such systems certainly provide for sharing knowledge claims, they provide no way of differentiating knowledge from mere information, so one cannot tell whether knowledge or information is being shared through them. In order for Best Practices systems to become KM interventions, they would need to incorporate meta-claims describing the track record of performance or at least the basis behind the Best Practice claims recorded in them. We have argued this at greater length elsewhere (Firestone and McElroy, 2003, Ch. 7).

Another popular "KM" intervention is the Enterprise Information/Knowledge Portal. One of us has distinguished these two types of portals sharply since early in 1999 (Firestone, 1999). But as the terms are used by most in KM, whether an application is called one or the other seems to be unrelated to any systematic difference in the interventions being discussed. In any event, portal tools and interventions, in spite of the early characterization of portals as "KM's killer app"

(Roberts-Witt, 1999), with perhaps a few exceptions for custom portal applications, fail to meet the above criteria for KM interventions (Firestone, 2003a, chs. 13-17).

Portals, like Best Practices systems, don't provide a way of distinguishing information from knowledge. As a consequence, any support they provide for integration functions such as broadcasting, sharing, teaching (through e-learning applications), and search and retrieval, is restricted to information, rather than knowledge, integration. Nor do portals generally provide targeted support for problem recognition, or for individual and group learning, or for knowledge claim evaluation. Nor do they provide targeted support for any of the KM activities distinguished in criterion 7 above.

There are the remaining possibilities that portal applications provide the required support for information acquisition and for knowledge claim formulation. But in the area of information acquisition, portal applications have shortcomings in the extent to which they support search results that are specifically relevant to problems. Though search technology has improved substantially since portals originated in 1998, it is widely recognized that they do not provide results that are sufficiently targeted on problems without a great deal of continuous interaction between humans and the portal.

Moving to knowledge claim formulation, many portal interventions focused on content management or collaborative capabilities do not provide support for idea management, semantic networking, formal modeling, simulation, or other techniques supporting alternative formulations. However, portals with strong structured data analysis/On-line Analytical Processing/Business Intelligence capabilities support knowledge claim formulation including the specification of alternative claims. These types of portals support knowledge processing and therefore interventions that deploy such portals are, indeed, KM interventions.

In brief, while portals provide a wide range of generalized support for information processing and management, portals focused on content management provide little specific support for knowledge processing as outlined in the criteria mentioned earlier. It is not impossible for portals to provide support in many of these areas, and hence for KM interventions based on portals to enhance knowledge processing. All it requires is that portal interventions incorporate portlets targeted at enhancing KLC functions. And, in fact, portals that support structured data analysis already provide support for knowledge claim formulation.

But portal studies (Firestone, 2003a, Collins, 2003, Terra and Gordon, 2003) show that they mostly focus on Content Management, Collaboration, Document Management, Publication, CRM, imprecise searching, publication, taxonomy development (a narrow type of knowledge production), and other forms of organizational support, that are not directly related to knowledge processing. Thus, many portal interventions are not KM interventions, and to determine which ones are requires analysis of the details of the portal application involved.

Turning to some examples from the area of social techniques for KM interventions we have listed, we think it is also the case that Communities of Practice (CoPs), Storytelling, and Social Network Analysis-based interventions *may or may not be KM interventions*, depending on the details of the specific intervention that is planned and implemented. Since CoP-based interventions are among the favorite initiatives of Knowledge Managers, we begin by asking the question, when is a CoP intervention not a KM intervention?

If the CoP intervention is aimed at enhancing knowledge sharing, but fails to provide a way of distinguishing CoP-produced Knowledge from CoP-produced information, then, we claim, it is not a KM intervention but an Information Management (IM) intervention. How widespread are such CoP-based interventions? While we have no data on this point, we believe that most CoP interventions are intended to enhance knowledge sharing but do not provide a way of distinguishing knowledge from information, and therefore that most are not KM interventions at all.

KM interventions that attempt to introduce the use of storytelling as a technique of knowledge sharing, share with CoP interventions the difficulty that they don't help to distinguish knowledge from information in what is shared. Stories are not automatically knowledge because humans tell them.

On the other hand, they are, automatically, a way of expressing knowledge claims, so that interventions enhancing the capacity to express knowledge claims in the context or form of stories may be viewed as KM interventions, assuming that they also enhance the capability to express alternative knowledge claims. In addition, interventions that enhance the storytelling capabilities of Knowledge Managers may be viewed as KM interventions, since they enhance both the leadership and knowledge claim formulation capabilities of Knowledge Managers.

A technique experiencing increasing popularity this year is Social Network Analysis (Cross and Parker, 2004), and one well-known KM blogger (Pollard, 2004) has even suggested that KM be re-invented as "social network enablement," meaning that KM interventions would aim at enhancing opportunities for social networks to form and thrive. Social Network Analysis (SNA) is clearly an analytic technique that can help generate knowledge claims about social networks, so interventions whose aim is to provide IT tools for performing SNA, or training in SNA, are certainly narrow-scope KM interventions since they enhance knowledge claim formulation including generating alternative social network models.

But Social Network Enablement as a management intervention is aimed directly at enhancing social network formation and maintenance and not at any knowledge process per se. Therefore, it cannot be a KM intervention technique, generally speaking, except when it is used to build KM-level external relationships, or as an aid in CoP or team-building interventions that are aimed at enhancing KM processes or various knowledge sub-processes in the KLC.

We hope the foregoing discussion of Best Practices Systems, EIPs, CoPs, Storytelling, and Social Network Enablement makes clear the following points. *Most interventions that have been viewed as KM interventions have not been instances of KM at all.* Nor is it possible, in many instances, to conclude that an intervention is a KM intervention based on the tool or social technique it uses. As the old saying goes, the devil is in the details, which, in turn, determine whether a particular intervention will fit one of the seven criteria we have specified earlier. In short, in many cases, where others think KM has been done frequently, our analysis implies that perhaps it has not been done. But having argued for that view, we now illustrate that KM both can be and has been done. Our illustration is the Partners HealthCare case to which we now turn.

The Partners HealthCare Case

In July, 2002, authors Tom Davenport and John Glaser published a case study in Harvard Business Review (Davenport and Glaser, 2002) involving a KM implementation at Partners HealthCare in Boston. Davenport is a KM researcher and consultant, and Glaser is the CIO at Partners.

The decision to invest in KM at Partners was largely driven by the cost of medical errors in healthcare, especially as reported by the Institute of Medicine (Kohn, Corrigan, and Donaldson, 1999) in 1998. According to IOM's report, more than a million injuries and as many as 98,000 deaths each year are attributable to medical errors. At Partners, medical errors, as measured by them in 1995, showed that "more than 5% of patients had adverse reactions to drugs while under medical care; 43% of those inpatient reactions were serious, life threatening, or fatal. Of the reactions that were preventable, more than half were caused by inappropriate drug prescriptions." (Davenport and Glaser, 2002, p. 5) Moreover, "A study of the six most common laboratory tests ordered by physicians in Brigham and Women's surgical intensive care unit found that almost half of the tests ordered were clinically unnecessary" (Davenport and Glaser, 2002, p. 6).

On the basis of these and other problems discovered at Partners, the decision to invest in KM was confined to the order-entry system, "because the system is central to physicians delivering good medical care. When doctors order tests, medications, or other forms of treatment, they're translating their judgments into actions. This is the moment when outside knowledge is most valuable. Without the system, doctors would have no easy way to access others' knowledge in real time." (Davenport and Glaser, 2002, p. 6)

Perhaps the most informative part of the Davenport and Glaser case was their description of how the KM system at Partners works. Here it is (Davenport and Glaser, 2002, p. 7):

"Here's how it works. Let's say Dr. Goldzer has a patient, Mrs. Johnson, and she has a serious infection. He decides to treat the infection with

ampicillin. As he logs on to the computer to order the drug, the system automatically checks her medical records for allergic reactions to any medications. She's never taken that particular medication, but she once had an allergic reaction to penicillin, a drug chemically similar to ampicillin. The computer brings that reaction to Goldzer's attention and asks if he wants to continue with the order. He asks the system what the allergic reaction was. It could have been something relatively minor, like a rash, or major, like going into shock. Mrs. Johnson's reaction was a rash. Goldzer decides to override the computer's recommendation and prescribe the original medication, judging that the positive benefit from the prescription outweighs the negative effects of a relatively minor and treatable rash. The system lets him do that, but it requires him to give a reason for overriding its recommendation."

Of central importance to the design of the integrated order-entry/KM system at Partners was the formation of centralized committees who were given the responsibility to "create and maintain the knowledge repository." (Davenport and Glaser, 2002, p. 8) Only "clinicians at the top of their game" (Davenport and Glaser, 2002, p. 8) were permitted to sit on these committees, and were given the authority "to identify, refine, and update the knowledge used in each [medical/clinical] domain." It was one of these committees of experts that was the source of the knowledge presented to Dr. Goldzer in the anecdote quoted above.

But despite the authoritative source of the knowledge presented to physicians at the time of order entry, Partners took a position of deference with respect to the decisions made by front line, practicing physicians in the hospital. "With high-end knowledge workers like physicians," they reasoned,

"it would be a mistake to remove them from the decision-making process; they might end up resenting or rejecting the system if it changed their role – and with good reason. Because over-reliance on computerized knowledge can easily lead to mistakes, Partners' system presents physicians with recommendations, not commands. The hope is that the physicians will combine their own knowledge with the system's." (Davenport and Glaser, 2002, p. 8)

As a result of the integrated order-entry/KM system at Partners, several benefits in the form of reduced medical errors were realized.

"Out of the 13,000 orders entered on an average day by physicians at Brigham and Women's, 386 are changed as a result of a computer suggestion. When medication allergies or conflict warnings are generated, a third to a half of the orders are canceled. The hospital's event-detection system generates more than 3,000 alerts per year; as a result of these alerts, treatments are changed 72% of the time – a sign that the hybrid human-computer system at Partners is working as it should." (Davenport and Glaser, 2002, pp. 8-9)

Also illustrative of the impact that KM had at Partners were the following results (Davenport and Glaser, 2002, p. 8):

- A controlled study of the system's impact on medication errors found that serious errors were reduced by 55%.
- When Partners experts established that a new drug was particularly beneficial for heart problems, orders for that drug increased from 12% to 81%.
- When the system began recommending that a cancer drug be given fewer times per day, the percent of orders entered for the lower frequency changed from 6% to 75%.
- When the system began to remind physicians that patients requiring bed rest also needed the blood thinner heparin, the frequency of prescriptions for that drug increased from 24% to 54%.

From this case, Davenport and Glaser concluded that the key to success in KM "is to bake specialized knowledge into the jobs of highly-skilled workers – to make the knowledge so readily accessible that it can't be avoided." (Davenport and Glaser, 2002, p. 6) They further conclude that:

"While there are several ways to bake knowledge into knowledge work, the most promising approach is to embed it into the technology that workers use to do their jobs. That ensures that knowledge management is no longer a separate activity requiring additional time and motivation.

We believe that this method could revolutionize knowledge management in the same way that just-in-time systems revolutionized inventory management – and by following much the same philosophy." (Davenport and Glaser, 2002, p. 6)

Analysis of the Partners HealthCare Case

Let's look at the Partners HealthCare case from the perspective of the conceptual frameworks used to evaluate other KM interventions. The three-tier framework suggests, first of all, that the intervention implementing the Partners system is a bona fide KM intervention, since its purposes appear to be to enhance: knowledge integration into DEC's and DOKBs, problem recognition and error elimination in DEC's, and knowledge production and the quality of knowledge in response to problems. On the other hand, the framework also suggests that the system is actually a knowledge processing system, rather than a "KM" system, as it is described by Davenport and Glaser. The knowledge processing system operates at two levels: the level of the individual doctor, and also the level of the organization.

- The Doctor's Level

One of the purposes of the Partners system was to reduce errors by upgrading knowledge at the point where doctors' make decisions to order tests, medications, or other forms of treatment. Knowledge at the point of decision was to be upgraded by way of the new system's ability to broadcast others' knowledge to the decision maker, *and also by the decision maker thereby using, or not, the shared knowledge to question his/her own decisions or actions.* In other words, from the point of view of our frameworks, the system is, in the first instance, *about eliminating or reducing errors in DEC's by increasing the frequency with which doctors question, critically evaluate, and recognize problems in the decisions they are contemplating.* The system is supposed to make doctors *look for problems* in their views, and if they find them, initiate problem solving (that is, KLCs) of their own, in the expectation that this will increase the quality of the beliefs that survive and inform their order entry decisions.

Thus, in terms of the DEC framework, when Dr. Goldzer uses his previous knowledge to decide to treat Mrs. Johnson's infection with ampicillin, he acts on the decision by ordering the drug. The system prepares to intervene in Dr. Goldzer's DEC between his action and the production of a result for him to monitor and evaluate. There were two options for the system in this situation. If it had not found any contra-indicating history (or other previous knowledge) related to Goldzer's order, his order would have been processed, and the results of Goldzer's DEC would have been the administration of ampicillin to Mrs. Johnson and its downstream effects.

The option applicable to Goldzer's actual situation, however, was that the knowledge claims in the system conflicted with his order, so the system intervened in Goldzer's DEC and brought Mrs. Johnson's previous allergic reaction to his attention by presenting him with a knowledge claim about that as the result of his decision. Thus, it *integrated the organization's knowledge into his DEC*, and forced him to *critically evaluate* his belief that the right thing to prescribe for Mrs. Johnson was ampicillin, against the knowledge claims it presented to him. In doing that, the system facilitated the possibility, or, if you like, increased the probability, that Dr. Goldzer *would question his decision, recognize a problem* with it, and then initiate a knowledge life cycle to solve it.

In fact, that is what he did, initiate a KLC, and specifically, from his individual perspective, an activity of information acquisition. When Dr. Goldzer learns that Mrs. Johnson's allergic reaction to penicillin was a rash, he uses that information and his judgment that "the positive benefit from the prescription outweighs the negative effects of a relatively minor and treatable rash," to falsify the computer's recommendation, the organization's knowledge, that he not prescribe ampicillin. He then ends his individual KLC and returns to the associated operational DEC, through which he again places his original order. Before he is allowed to proceed, however, he is prompted by the system to integrate into the

organizational DOKB knowledge claims and meta-claims explaining why he falsified (over-rode) the system's knowledge claims.

- *The Organizational Perspective*

When we look at the order entry system from an organizational perspective, we see knowledge production being performed by committees of experts. They evaluate what goes into the system, and the claims they approve receive the "imprimatur" of the organization as knowledge to be integrated into order entry DEC's when triggered by specific transactions. In terms of Dr. Goldzer's activities, one of the committees of experts was the source of the order-relevant knowledge presented to him in the description quoted above.

From the viewpoint of our frameworks, the committees are continuously processing knowledge claims in an effort to reduce or eliminate the errors in the DOKB, and thus to upgrade its quality over time. *The committees are designated authorities for knowledge production and knowledge claim evaluation* at the organizational level, directed at solving the problem of medical error reduction in order entry. They perform KLCs, evaluate, and select the knowledge claims that are formally designated as organizational knowledge, and that will be made available through the system for integration into the order entry DEC's.

The system however, works in such a way that the centralization of knowledge production in the committee is balanced by the participation of all physicians in knowledge claim formulation and evaluation in the context of their participation in the order entry system. Partners understood the need to maintain a distributed decision making system with respect to order entry, and, in addition, to reinforce *a distributed problem solving system* with respect to problems arising out of the order entry decision. Partners did this because it recognized *the fallibility of organizational knowledge* produced by the committees, the need to involve the doctors and their knowledge in solving problems and adding knowledge claims to the DOKB, and *the need to view system interventions in decisions made by the doctors, as acts of knowledge integration*, intended to strengthen monitoring and evaluation and problem recognition in the DEC, rather than knowledge imposition.

In the end, the Partners system is stronger because it is a distributed problem solving system, in which the committees, through the system, help the doctors to recognize that there are problems with some of their orders. But by sometimes insisting on their decisions and giving the committees feedback on their own reasons for doing so, the doctors, again through the system, are providing knowledge claims to the committees, as well as critical evaluations of the committees' recommendations (i.e., their knowledge claims) to them, in the form of the reasons they provide for over-riding such recommendations.

When the committees later review the doctors' claims, and decide whether to incorporate them into the system and/or modify their own previous recommendations, they are engaging in further knowledge claim evaluation and

in producing new knowledge at the organizational level. In this way, *the system links the individual level with the organizational level and makes the doctors participants, along with the committees, in organizational problem solving and knowledge production.* In this regard, the Partners system not only injects organizational knowledge at key decision points in the order entry process, but also integrates knowledge processing functionality at the same time and place in the form of knowledge claim formulation and knowledge claim evaluation.

On the one hand, the system broadcasts organizational knowledge to the physicians and supports further information retrieval as well, while on the other it engages them in various aspects of organizational knowledge processing. Of most significance here is the requirement that the physician, Dr. Goldzer, record his reasons for over-riding the expert committee's recommendations. There's a dialectical dimension to the system. This suggests a new metaphor or heuristic for KM. It's not just *push versus pull* anymore, it's *push or pull and pull back!*

Finally, Davenport and Glaser, in their account, characterize the Partners system as embedding KM into the business process, or "baking specialized knowledge into knowledge work." From the viewpoint of our frameworks, however, KM remains where it is, in the top-tier of the framework. The Partners system supports knowledge processing at both individual and organizational levels. Organizational knowledge is integrated into DEC's so that problems are surfaced. More KLCs occur, evaluating the knowledge of both doctors and the organization, and, lastly, errors are more likely to be reduced or eliminated.

Implications for KM Strategy and KM Programs

The Partners HealthCare case is a great illustration of how to go about a successful KM intervention that enhances knowledge processing at the levels of both the individual and the organization in such a way that the changes have an impact on business outcomes: in this case, lives saved and serious consequences of medical errors avoided. The case also leads us to suggest an extension of the pattern into a KM strategy that, we propose, is at once coherent and incremental.

The vision of the strategy is to gradually enhance knowledge processing in the enterprise in a manner that will add increasing value and create sustainable innovation over time. The end state of the strategy is attaining a form of organization called the Open Enterprise, which, theory suggests, is an environment providing maximal support for sustainable innovation, problem solving, and adaptation. The Open Enterprise, an extension of Karl Popper's ideas about the Open Society (Popper, 1945) to organizations, is open to:

- New problems recognized by any of its agents
- New ideas generated by any of its agents (knowledge claim formulation)

- Continuous criticism of previously generated ideas by any of its agents (knowledge claim evaluation)

The Open Enterprise is not democratic in decision making or in management. But it requires at least internal transparency and inclusiveness in distributed knowledge processing and problem-solving. Here are the steps in the strategy.

1. Use a formal KM methodology to implement the strategy.

There are many methodologies available that apply to KM tools and techniques, but there is very little in the literature offering a comprehensive KM program and project methodology. Perhaps the only alternative is K-STREAM™, a recent formulation of our own (Firestone and McElroy, 2004a, 2004b, KMCI, 2004).

2. Identify and Prioritize decisions (DECs), work flows, or business processes according to risk.

In formulating a KM strategy and an associated program, one needs to systematically specify DECs and, where necessary, work flows, or business processes that can produce highly negative business outcomes if errors are made. In the Partners case, the organization identified a decision, the order entry decision, involving high risk for the organization. That decision was the source of costly medical errors, and having a favorable impact on it was likely to produce a lot of social credit for the KM function at Partners. Identification of high risk DECs should be followed by prioritization of them according to risk, taking into account, ease and expense of intervention, and likelihood of success. It should also be understood that whereas high-risk decisions are the logical place to start in this approach, decision-oriented KM interventions do not end there. All decisions in organizations are subject to enhancement, and the opportunity to improve performance in more general terms exists across the board. Demonstrating success for high-risk decisions first, however, is appropriate, both for purposes of having impact early in KM programs and also for building confidence and justification for further investment in KM.

3. Select DECs, work flows, or business processes as targets for KM interventions according to priority and develop the business case.

There is no indication that Partners selected their KM intervention in the order entry decision from a set of alternatives. But an essential step in developing a long-term strategy is to use the outcome of step 1 to perform such a selection, and in doing so to develop the outlines of a KM program designed to enhance knowledge processing.

4. If you can, make interventions that embed new knowledge processing functionality within existing IT-based business applications supporting DECs, work flows or business processes.

This follows the pattern of the Partners case. It assumes that you can find high risk DEC's already supported by existing IT applications to use as the objects of intervention. Intervening in existing applications is preferable to introducing entirely new applications, because people already depend on these applications as part of their job, and are likely to continue to use the enhanced system. Note here, however, that the integration of knowledge processing functions at key decision points need not necessarily take the form of IT implementations. Process or procedural changes can also be made, the effects of which will cause key decisions, and the knowledge claims behind them, to be tested by others before being put into action.

5. Make sure the new functionality added to the IT business application, or process, presents competing organizational knowledge claims to those expressed or implied in a DEC outcome.

This is needed to encourage questioning of previous individual-level knowledge in DEC's, which, in turn, can encourage increased problem recognition, individual KLCs, and error reduction in key decisions. Knowledge processing systems resulting from such KM interventions will help decision makers to "look for trouble," recognize problems, and initiate KLCs, and in the process will bring inclusiveness to problem recognition and problem solving related to the high risk area which is the target of the intervention.

6. When competing knowledge claims introduced in knowledge integration are over-ridden by a decision maker, new IT application (or process) functionality should require that the superceding knowledge claims and meta-claims be added to the DOKB by the decision maker.

This is essential to accumulate a track record of knowledge claim performance in the DOKB. Soliciting knowledge claims and meta-claims in this way opens up knowledge processing to new ideas and to distributed knowledge claim evaluation. Thus, it moves the organization closer to the Open Enterprise by including the decision maker (knowledge worker) in knowledge processing, and in knowledge production specifically.

7. Once the first intervention is completed, continue implementing the KM program, project by project, according to the priority established earlier.

Following the strategy will strengthen the ability to: recognize problems in area after area, initiate KLCs, produce distributed problem solving, and increase adaptiveness. And in the process it will move the organization closer to the Open Enterprise, problem area by problem area, through creating transparency, inclusiveness and other characteristics of the Open Enterprise in each case. That progress should be tracked and measured, since the closer the organization gets to the Open Enterprise, the more it will exhibit adaptiveness and sustainable innovation.

Summary and Conclusions

KM as a field has been characterized by a great deal of confusion about its conceptual foundations and scope. As a result, practitioners have tended to view KM interventions as those that have been given that name by themselves or others who claim to be practitioners. In this paper, we have suggested that continuing that practice is destructive to KM as a discipline, because it prevents coherent evaluations of KM's track record. Moreover we have (a) offered a framework and set of criteria based on it for deciding whether claimed interventions are bona fide instances of KM, and (b) illustrated the use of that framework in critical evaluation of typical "KM" interventions, including extensive discussion of an unambiguous case where KM has been done.

This case, the well-known Partners HealthCare project, was also shown to illustrate a pattern of intervention that can serve as the basis of a long-term systematic strategy for implementing KM in the enterprise. The strategy is risk-based. It is one that can deliver concrete, incremental solutions and benefits to the enterprise by creating *quality-control systems for knowledge-in-use* as a support for distributed decision making and knowledge processing. In the long run, it can transform the enterprise into an organizational form that we call the Open Enterprise, and thereby support sustainable innovation and help solve the general problem of organizational adaptiveness and performance.

References

- Ackoff, R. (1970), *A Concept of Corporate Planning*, New York, NY: Wiley-Interscience.
- Allee, V. (2003), *The Future of Knowledge*, Burlington, MA: Butterworth-Heinemann, 2003.
- Argyris, C. (1993). *Knowledge for Action*, San Francisco, CA: Jossey-Bass.
- Argyris, C. and Schön, D. (1974) *Theory in Practice: Increasing Professional Effectiveness*, San Francisco: Jossey-Bass.
- Bickhard, M. (1999), "Information and Representation in Autonomous Agents," available at: <http://www.lehigh.edu/~mhb0/infrepautagents.html>
- Bickhard, M. (2000), "The Dynamic Emergence of Representation," available at: <http://www.lehigh.edu/~mhb0/dynemergrep.html>
- Campbell, D. (1974), "'Downward Causation' in Hierarchically Organized Biological Systems", in F. Ayala, and T. Dobzhansky, (eds.), *Studies in the Philosophy of Biology*, London, UK: MacMillan.
- Cavaleri, S. and Reed, F. (2000), "Designing Knowledge Generating Processes," *Knowledge and Innovation: Journal of the KMCI*, Vol. 1, no. 1, pp. 109-131.

Cavaleri, S. and Reed, F. (2001), "Organizational Inquiry: The Search for Effective Knowledge," *Knowledge and Innovation: Journal of the KMCI*, Vol. 1, no. 3, pp. 27-54.

Collins, H. (2003), *Enterprise Knowledge Portals*, New York, NY: AMACOM.

Cross, R. and Parker, A. (2004), *The Hidden Power of Social Networks*, Boston, MA: HBSP, 2004.

Davenport, T. and Glaser, J. (2002), "Just-in-Time Delivery Comes to Knowledge Management," *Harvard Business Review*, July, 2002, pp. 107-111.

Denning, S. (2001), *The Springboard*, Boston, MA: KMCI Press/Butterworth Heinemann.

Firestone, J. M. (1999). "Enterprise information Portals and Enterprise Knowledge Portals," *DKMS Brief*, **8**, *Executive Information Systems, Inc.*, Wilmington, DE, March 20, 1999. Available at http://www.dkms.com/White_Papers.htm

Firestone, J. M. (2000) "Knowledge Management: A Framework for Analysis and Measurement," *White Paper No. 17*, *Executive Information Systems, Inc.*, Wilmington, DE, October 1, 2000, Available at: http://www.dkms.com/White_Papers.htm.

Firestone, J. M. (2000a), "Enterprise Knowledge Portals: What They Are and What They Do," *Knowledge and Innovation: Journal of the KMCI*, Vol. 1, no. 1, pp. 85-108. Available at: http://www.dkms.com/White_Papers.htm.

Firestone, J. M. (2003), "Minding the (Knowledge) Gap," *Knowledge Management Magazine*, Vol. 6, no. 8 (April, 2003), pp. 20-24.

Firestone, J. (2003a), *Enterprise Information Portals and Knowledge Management*, Burlington, MA: KMCI Press/Butterworth Heinemann.

Firestone, J. M. (2003b), "How Knowledge Management Can Help Identify and Bridge Knowledge Gaps," *An EIS Professional Paper*, Wilmington, DE: Executive Information Systems, Inc., 2003, available at: <http://www.dkms.com/professionalpapers.htm>.

Firestone, J. and McElroy, M. (2003), *Key Issues in the New Knowledge Management*, Burlington, MA: KMCI Press/Butterworth Heinemann.

Firestone, J. and McElroy, M. (2003a), "The New Knowledge Management," *Knowledge Management Magazine*, Vol. 6, no. 10 (June, 2003), pp. 12-16.

Firestone, J. and McElroy, M. (2003b), *Excerpt #1 From The Open Enterprise: Building Business Architectures for Openness and Sustainable Innovation*, (with Mark W. McElroy) (Hartland Four Corners, VT: KMCI Online Press, 2003), available at: <http://www.dkms.com>, <http://www.macroinnovation.com>, and <http://www.kmci.org>

Firestone, J. and McElroy, M. (2004), "Viewpoint: Organizational Learning and Knowledge Management: the Relationship," *The Learning Organization*, Vol. 11, no. 2, pp. 177-184.

Firestone, J. and McElroy, M. (2004a), "K-STREAM™ and the New Knowledge Management?" *All Life is Problem Solving*, May 12, 2004, available at: <http://radio.weblogs.com/0135950/>

Firestone, J. and McElroy, M. (2004b), "Are Their Core Tools and Techniques of Knowledge Management?" *All Life is Problem Solving*, May 12, 2004, available at: <http://radio.weblogs.com/0135950/>

Gell-Mann, M. (1994), *The Quark and the Jaguar*, New York, NY: W. H. Freeman.

Haeckel, S. H., (1999). *Adaptive Enterprise*, Boston, MA: Harvard Business School Press.

Hall, W. (2005), "The Biological Nature of Knowledge in Learning Organizations," *The Learning Organization*, This issue.

Holland, J. H. (1995). *Hidden Order*. Reading, MA: Addison-Wesley.

Holland, J. H. (1998). *Emergence*. Reading, MA: Addison-Wesley.

Isaacs, David (1999) "Knowledge Café Presentation," Enterprise Intelligence Conference, Dec. 7, 1999, Lake Buena Vista, FL

Juarrero, A. (1999), *Dynamics in Action*, Cambridge, MA: MIT Press.

Kauffman, S. (1995). *At Home in the Universe*. New York, NY: Oxford University Press.

KMCI (2004), "CKIM Certificate Program," available at: <http://www.kmci.org/ckim-certificate.html>

Kohn, L., Corrigan, J. and Donaldson, M. (eds.) (1999), *To Err is Human: Building a Safer Health System, Committee on Quality of Health Care in America, Institute of Medicine*, Washington, D.C.: National Academy Press, 1999.

Kolb, D. (1984), *Experiential Learning*, Englewood, Cliffs, NJ: Prentice-Hall.

Kolb, D. and Fry, R. (1975), "Toward and Applied Theory of Experiential Learning," C. Cooper (ed.), *Theories of Group Process*, London: John Wiley and Sons.

McElroy, M. W. (1999). "The Second Generation of KM," *Knowledge Management* (October, 1999), pp. 86-88.

McElroy, M. W. (2000a), "The New Knowledge Management," *Knowledge and Innovation: Journal of the KMCI*, Vol. 1, no. 1 (October 15, 2000), pp. 43-67.

McElroy, M. W. (2003). *The New Knowledge Management: Complexity, Learning, and Sustainable Innovation*, Burlington, MA: KMCI Press/Butterworth-Heinemann.

Pollard, D. (2004), "Social Networking, Social Software and The Future of Knowledge Management," *How to Save the World*, available at: <http://blogs.salon.com/0002007/2003/05/28.html#a251>

Popper, K. R. (1945), *The Open Society and Its Enemies*, London, UK: Routledge & Sons.

Popper, K. R. (1972), *Objective Knowledge*, London, England: Oxford University Press.

Popper, K. R. (1978), "Three Worlds: The Tanner Lecture on Human Values," available at: <http://www.tannerlectures.utah.edu/lectures/popper80.pdf>

Popper, K. R. (1987), "Natural Selection and the Emergence of Mind," in Gerard Radnitzky and W. W. Bartley, III, (eds.), *Evolutionary Epistemology, Rationality, and the Sociology of Knowledge*, La Salle, IL: Open Court.

Popper, K. R. (1994). *Knowledge and the Body-Mind Problem*, (edited by Mark A. Notturmo), London, UK: Routledge.

Popper, K. R. (1999), *All Life is Problem Solving*, London, UK: Routledge.

Popper, K. R. and Eccles, J. C. (1977). *The Self and Its Brain*, Berlin, Germany: Springer.

Roberts-Witt, S. L. "Making Sense of Portal Pandemonium," *KM Magazine* (July, 1999), pp. 37-53.

Senge, P. (1990). *The Fifth Discipline*, New York, NY: Currency Doubleday.

Terra, J. and Gordon, C. (2003), *Realizing the Promise of Corporate Portals*, Burlington, MA: KMCI Press/Butterworth-Heinemann

Wenger, E. (1998). *Communities of Practice: Learning, Meaning, and Identity*, New York, NY: Cambridge University Press.

About the Authors

Joseph M. Firestone and Mark W. McElroy are both very well known in the field of KM and have been collaborating with one another since 1998. Each has his own organizational affiliation, however, as indicated at the start of the paper, with Firestone's firm, Executive Information Systems, Inc., to be found at www.dkms.com, and McElroy's organization, Center for Sustainable Innovation, located at www.sustainableinnovation.org.

About This Paper

This paper was written for publication in a special issue of *The Learning Organization* journal co-edited by Firestone and McElroy, the focus of which was the rhetorical question of *Has KM Been Done?*.