

Aus: Joseph M. Firestone, PH.D. Mark W. McElroy:

Key Issues in the New Knowledge Management

Chapter 11 CONCLUSION

VISION OF THE NEW KNOWLEDGE MANAGEMENT

The new Knowledge Management is more than just the second generation of KM-it is a new science, a social science. It is a science because its theory and practice are testable and falsifiable, yet systematic and formal, and it is a social science because of the nature of its focus on how to get human social systems to learn and adapt as best they can. The new KM can also be seen as an application of a much broader "new science": complexity science. Thus, it is a new application of Complex Adaptive Systems (CAS) theory, a vision of how living systems adapt to changes in their environment by evolving their knowledge. And because of the human social setting of its focus, the new KM also relies on organizational learning (OL) (as well as social psychological, sociological, and political) theory to round out its views. Not only do individuals learn, but the groups and organizations of which they are a part also learn. In this way, organizational learning theory, along with other social sciences that contribute to KM, help us put a human face on CAS theory, thereby accounting for the unique blend we see between OL, other social science disciplines, and CAS in the new KM.

Despite their contributions to KM, organizational learning, other social sciences, and complexity theories still need to be supplemented in one very important area: epistemology. This is (or was) a problem. After all, we can't have a discipline called "Knowledge Management" that fails to adequately define its own vocabulary, or to ground itself on a theory of knowledge. For that we turned to Karl Popper (1963, 1972, and 1994; Popper and Eccles 1977). Popper's notion of the three "worlds" of existence is comprehensive and convincing. Not only does it account for knowledge held in minds (beliefs), it also provides for so-called "objective" knowledge-knowledge held in linguistic expressions, such as speech, documents, and other material recordings. In organizations, we find both "subjective" knowledge held in minds and "objective" knowledge held, or expressed in linguistic forms.

Popper's epistemology is also highly compatible with organizational learning and complexity theory. To the former (as we have shown in Chapter 2), he adds the dynamics of knowledge claim formulation and trial-and-error as the means by which "double-loop learning" occurs; to the latter, a focus on the elimination of errors in our thinking and how error elimination can enhance our capacity to adapt. In addition, his theory of the emergent evolution of the three ontological worlds emphasizes the adaptive significance of World 2 (subjective)

and World 3 (objective) knowledge in human evolution (see Chapter 1). That is, according to Popper (1972, 1994; Popper and Eccles 1977), *both types of knowledge are for adaptation, and the ability to produce both has been selected for by the environment. This is the significance of knowledge in human evolution.*

Indeed, Popper's views give managers, and other members of adaptive systems who wish to adapt, a better and more effective place to go in terms of how to approach the production of their own knowledge than do alternative theories of knowledge. That is, his views make it clear *that knowledge production by them IS a legitimate and natural process that they must engage in if they and their organizations are to adapt.* Similarly, Popper's epistemology gives new meaning to the practice of KM-the new KM, that *is-by orienting organizations toward seeking out and eliminating the errors in their knowledge claims.* Institutionalizing related processes on an enterprise-wide basis is an important aspect of the mission and purpose of the new KM, because of the centrality of error elimination to adaptation.

The old KM-still widely practiced in many circles-is not entirely free of theoretical foundations itself. And it may even qualify, for some, as a science. But on both fronts, it is quite different from the new KM and considerably narrower. While mostly unspoken, we could say, for example, that the theory of first-generation KM goes something like this: *that we can assume the existence of pre-existing knowledge and that KM is about getting that "knowledge" (actually "information") to the right people at the right time.* The scope of first-generation KM's interests, therefore, is (and has been) the delivery of existing information. We say "Information" and not "knowledge" because no persuasive distinction between the two has ever been made in the conventional practice of (old) KM, and the term "information" is used in the KM literature at least as often, it seems, as the term "knowledge" is.

Thus, the old KM is largely predicated on the view that valuable knowledge simply exists and that the purpose of KM, therefore, is to enhance its delivery to people who need it. But since no meaningful distinction between knowledge and information has ever been made in the practice of first-generation KM, we feel safe in concluding that the old KM has not been about "knowledge" management at all-it's been about "information" management (see Chapter 3).

Further, the assumption that knowledge (information) simply exists and that KM's role should merely be to aid in its capture and delivery is also unfounded, or at least arbitrary. While it may be convenient for some to ignore the means by which knowledge is produced-and not just *shared* or delivered-that still leaves the rest of us, who actually from time to time have a need to learn, innovate, and adapt, dissatisfied. More important still, *the essence of life, and of individuals and human organizations, is . is adaptation to changes in the environment. The old KM is not about such adaptation, and therefore it is basically incrementalist in orientation . on and not focused on the most*

essential of organizational functions. Compared to the new KM, then, it is unimportant, a mere diversion from the fundamental questions of organizational existence and development.

The new KM, then, is aimed at enhancing the whole cycle of knowledge production and integration, not just the easier and more convenient tasks of rounding up documents and other artifacts for access and distribution in computer-based repositories. Its scope of attention is, therefore, considerably broader and deeper than its first-generation cousin, as are its value propositions. Moreover, the distinction it makes between knowledge and information takes it (and us) into utterly new areas of insight and opportunity. There is a difference between knowledge and information, and the new KM opens our eyes to it for the first time in the practice of KM.

Perhaps the most dramatic example of this new perspective is the connection we have drawn between KM and corporate malfeasance. Turning to Popper, we can interpret such behaviors as *error-ridden knowledge in practice*, employed by unscrupulous corporate executives, whose knowledge in use is not only faulty, but whose knowledge in use somehow managed to get a great deal farther along than it should have. Thus, such knowledge in use is riddled with errors, and the processes in place that produced and validated it (i.e., the knowledge processing processes) are dysfunctional, as well.

By focusing on enhancing the quality and integrity of organizational knowledge processing, the new KM can play an extremely valuable role in helping to reduce corporate malfeasance. Indeed, by calling attention to the *dimensions of errors* in our knowledge and to the importance of the processes used to produce and evaluate it, we can have a material impact on improving the quality of the business processing outcomes that follow. But since the old, first-generation KM makes no distinction between knowledge processing and business processing-much less between knowledge and information in terms of errors, validity, or anything else-no such role for KM could possibly come out of it. Its sole purpose is to enhance the delivery of information.

Deep in the heart of the new KM are two Principal claims that perhaps underlie its practice more than anything else. The first, taken from Popper (1959) and from Peirce (1955), is that *there can be no certainty of truth and therefore no certain knowledge. All human knowledge is fallible.* We can, however, endeavor to get closer to the truth, and the way we can is by seeking to find and eliminate the errors in our thinking. This approach to learning is what Popper (1963) referred to as fallibilism.

The second insight follows closely on the heels of the first in the sense that *what people, in fact, seem to do as they go about the process of making decisions and learning is problem recognition, solution formulation, and error elimination through trial and error.* In this book (Chapter 2) and in other places we have described how trial and error works in the Decision

Execution Cycle (DEC) (Firestone 2000, Firestone 2003), and also in knowledge claim formulation and evaluation in the KLC.

When we combine these two ideas, what we get is a fundamental dynamic found in all humans: the need to solve problems by reaching tentative solutions and then eliminating the errors in them through trial-and-error processes of testing and evaluation as a precursor to decisions and action. Thus, the difference between knowledge and information is that while both consist of claims about the world, beauty, and truth, only knowledge consists of claims that have survived our tests and evaluations of whether they are in error. This doesn't make them necessarily true, but it does make them preferable to us as a basis for decisions and action, and that is why we pay so much attention to the idea of metaclaims and the role that they play in knowledge processing in the new KM.

Metaclaims are the claims about claims that we rely on (and that we also produce) as we seek to weigh the closeness to the truth of one claim against another by deciding which of the claims has survived our tests. Trying to distinguish surviving knowledge claims from false ones in the absence of metaclaims is a little bit like trying to find "middle C" on a piano keyboard whose black keys are missing. Without them, there's no way to distinguish one white key from any other.

Knowledge, then (but not truth), is entirely relative to the results of our tests and evaluations. What passes for knowledge to us (Firestone and McElroy) may amount to false or untested claims to you because your record of testing and evaluation may differ from ours. Moreover, some knowledge claims may be closer to the truth than others, and some may even be true, but in the end all human knowledge is fallible, and so we can never know with certainty which of our beliefs or claims are actually closest to the truth, or even coincident with it.

The best we can do is to demonstrate the absence of errors in our knowledge to the best of our ability by subjecting it to continuous criticism, testing, and evaluation. People and organizations that do this well stand a better chance of adapting and surviving in the face of change because the quality of their knowledge claims will be higher, and so their performance using them will also be superior.

THE NEW KNOWLEDGE MANAGEMENT LANDSCAPE

In this book we have tried to illustrate the many ways in which The New Knowledge Management (TNKM) differs from its firstgeneration cousin using key issues as a vehicle for doing so. We have covered:

- The difficulties KM has had in arriving at a cogent account of knowledge;
- The problem of explaining the origin of the distinction between business processing and knowledge processing and of the origin of the KLC;
- The distinction between information management and KM and the specification of a KM framework;
- The theory of change in generations of KM;
- The nature of knowledge claim evaluation and a framework for analyzing it and performing it;
- The significance of the KLC and some of its many applications;
- The meaning of best practices and some requirements for viable best practices systems;
- The problem of KM strategy, its independence of operational management, and the question of its independent role and status in the adaptive organization;
- The problem of the role of culture in KM and the overestimation of this role in KM to date;
- The progress of the field of intellectual capital and its failure to take account of Social Innovation Capital (SIC).

There are, however, many other aspects of the landscape of The New 'M (TNKM) that, for lack of space, we have left for examination a future volume. These additional key issues on the KM landscape include:

- The SECI Model (Nonaka and Takeuchi 1995) and the need for an expanded knowledge conversion framework;
- Is KM ready for standards?
- Portal progress and KM;
- A framework for analysis and evaluation of KM software;
- TNKM metrics;
- The open enterprise and sustainable innovation;
- TNKM and the terrorism crisis;
- The role of credit assignment systems;
- Communities of practice versus communities of inquiry;
- KM methodology: linear life cycle versus adaptive "tool box" approaches;
- KM methodology: knowledge audits versus knowledge process audits;
- KM methodology: the role of the KLC in KM methodology.

We will come back to these in our discussion of the future of TNKM. Let us now review, in summary form, a few of the most important ideas that arose from our discussion of the issues covered in the preceding chapters, while keeping in mind that our discussion neglected or only touched upon many important issues in the KM landscape.

More on Defining Knowledge

In addition to the analysis already offered in our section on the vision of TNKM, here are a number of other conclusions on the nature of knowledge in TNKM. In the old KM, the terms "knowledge" and "information" are used synonymously and without distinction, but in the new KM, as indicated in Chapter 1 and many other chapters in this book and the last section, we address the issue of making the distinction head-on. We do so by calling attention to the general ideas of knowledge claim evaluation and error elimination and to the testing and evaluation of knowledge claims. Closely associated with these ideas is the idea of claims and metaclaims, and the crucial role that metaclaims play in helping us to pick and choose from among knowledge claims based on how well they stand up to our tests-our attempts to falsify them.

In addition to knowledge in the form of claims, *knowledge also exists in the form of unexpressed beliefs and belief predispositions. This kind of knowledge,*

unlike claims, is held in minds-it is mental knowledge, not artifact-based or cultural knowledge. Since one person can never directly know what another one believes, mental knowledge is personal and subjective. In order to share knowledge and make it accessible to others, it must be expressed in some way-linguistically-thereby producing "objective" knowledge in the form of a claim. But this claim is not identical to the unexpressed belief or belief predisposition. There is an irreducible epistemic gap between what we believe and what we say.

In the organizational setting, we find both kinds of knowledge: subjective knowledge (what we believe) and objective knowledge (what we say). Further, both kinds of knowledge can be held, or expressed, in different ways, such as tacitly, implicitly, and explicitly. Indeed, in the new KM we expand greatly beyond the SECI conversion model (see Figure 11. 1) put forth by Nonaka and Takeuchi (1995), first by asserting the existence of not just tacit and explicit knowledge, but implicit knowledge, as well; and, second, by applying all three forms to both subjective and objective knowledge. In a future work of ours, we will return to this idea and elaborate on it in more detail. Its significance, however, is that *the SECI model greatly oversimplifies knowledge conversion, and therefore KM programs based upon it neglect many important forms of knowledge conversion.* We will say a little bit more about the SECI model below.

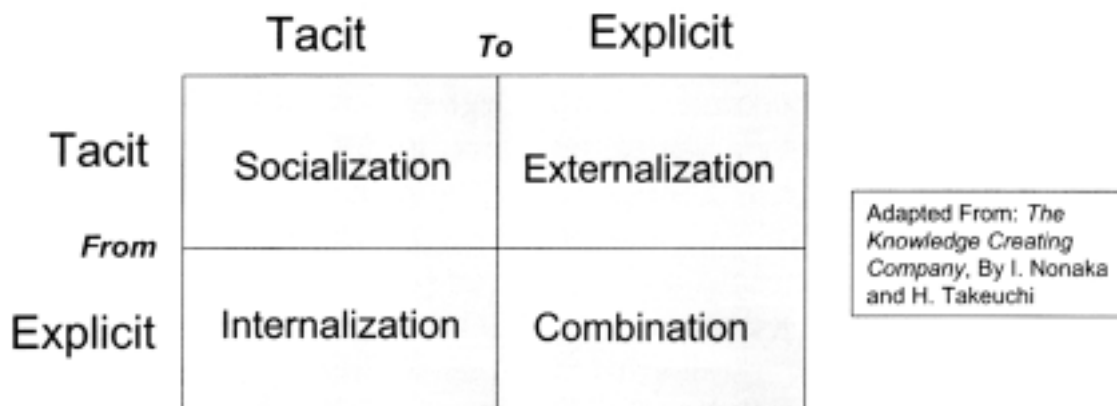


Figure 11. 1
The SECI Model

What, then, are the KM implications of the claims and metaclaims perspective in the new KM? Most important, this perspective tells us that not only is there a difference between knowledge and information, but that artifact-based knowledge is a particular kind of information. Information is not a subset of knowledge; knowledge is a subset of information! What kind of subset?

In the case of objective knowledge, it is a linguistic subset consisting of claims that have survived our tests and evaluations attempting to falsify it and whose record of testing and evaluation is accessible to us. Thus, we can say that

a knowledge processing system (be it social or technological in character) should not be regarded by us as a bona fide knowledge processing system unless it reveals and reports both knowledge claims and their metaclaims. Similarly, a KM strategy or intervention is not about KM at all unless it concerns the implementation or management of systems (social, technological or otherwise) that deal directly with the distinction between claims and metaclaims.

Here, we point out that *according to the new KM, most instantiations of KM in the past have not been about KM at all.* Perhaps this, in large part, accounts for the degree to which KM has been received with such ambivalence, reticence, and downright distrust to date. It never really *was* about *knowledge* management at all.

The Origin of the KLC

In Chapter 2, we provided a step-by-step development of the origin of the KLC in the Decision Execution Cycles (DECS) (sometimes referred to as Organizational Learning Cycles, or OLCS) that account for activity in human social systems and that, when interrelated by goals and objectives, form business processes. *We offered a theory of how DECs (see Figure 11.2), when characterized by perceived inadequate belief knowledge, can provide a context for recognition of a gap between what the decision maker knows and what he or she needs to know to make a decision. This epistemic gap is what we mean by a "problem." Problems, in turn, arouse a learning incentive. A system that motivates performance in a set of interrelated DECs aimed at closing the epistemic gap or solving the problem.*

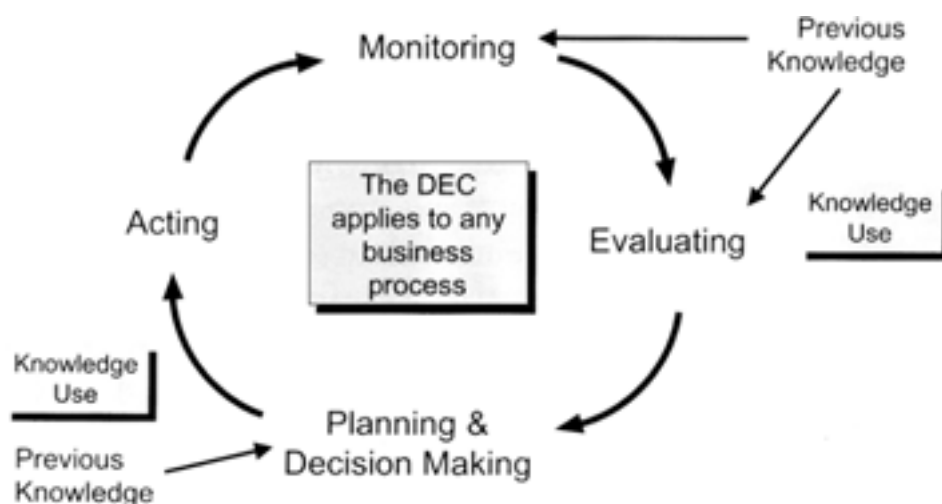


Figure 11.2
The Decision Execution Cycle (DEC)

This set of interrelated DEC's is, in fact, the KLC pattern, where DEC's are structured into information acquisition, individual and group learning, knowledge claim formulation, knowledge claim evaluation, broadcasting, searching / retrieving, sharing, and teaching. These are the subprocesses of the KLC (see Figure 11.3)-processes of knowledge production and integration that produce the distributed organizational knowledge base, all of the mental and artifact-based knowledge that originates in the enterprise.

We have seen (in Chapter 2) that the alternation between DEC's I. in business processing and DEC's in knowledge processing is basic to adaptation and grounded in human psychology, both at the individual and group levels. It is an alternation between different types of motivation, and this alternation is the foundation of a distinction between business processing and knowledge processing and between the latter and knowledge management. This last distinction is the basis of knowledge management as a distinct process and discipline. Without it there can be no knowledge management.

Knowledge Process Management and Information Management

A second key insight, now plainly obvious to many of us thanks to the new KM (see Chapters 3 and 4), is that managing "knowledge" as an outcome—even with our clarified understanding of what it is—has been a mistaken notion all along. *KM is not just about knowledge (claims and metaclaims) outcomes management; rather, it's about knowledge process management, as well.* Once we recognize and come to terms with the fact that knowledge is produced in organizations by way of a social process, it's a short hop from there to realizing that *we can only manage the outcomes of such a process through managing the process itself.*

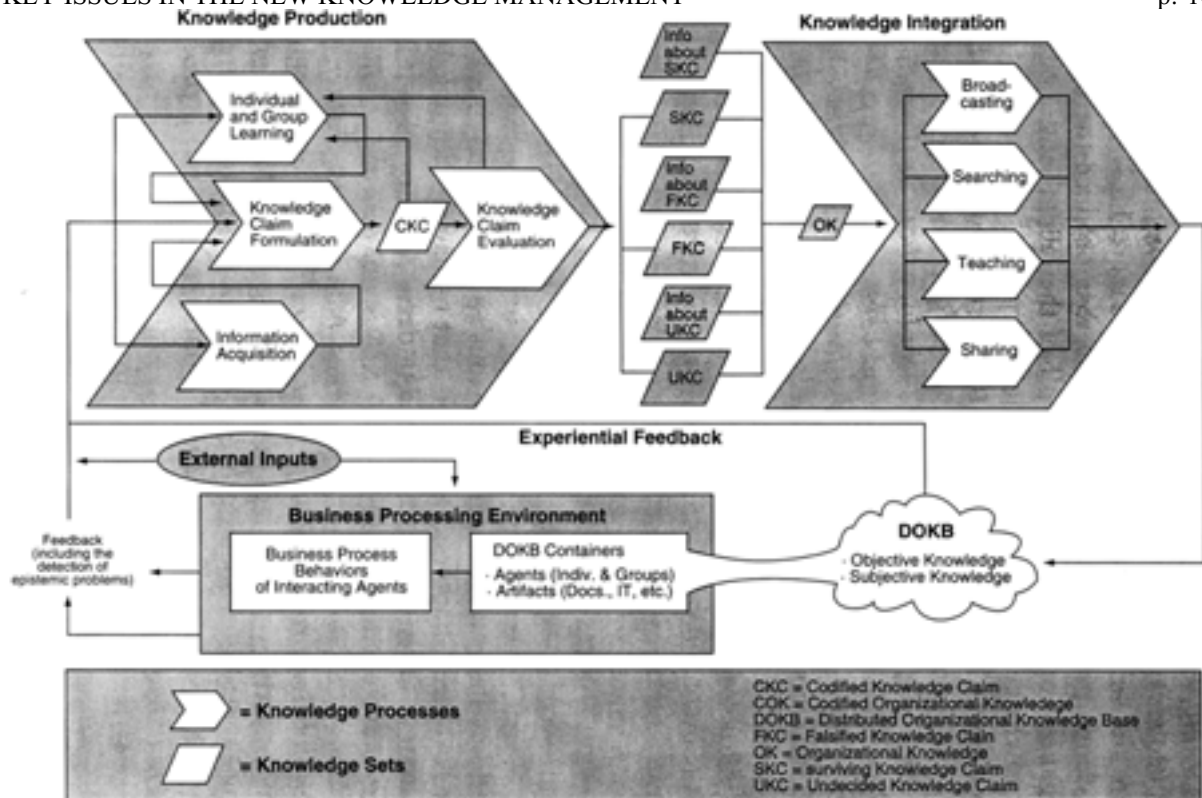


Figure 11.3
The Knowledge Life Cycle (KLC)

Here is where the new KM reveals another crucial weakness of the old KM—the premise that valuable knowledge already (or simply) exists and that the proper role of KM is merely to find it, codify it, and help make it more easily accessible to people who need it. But clearly this is not the entire story. Knowledge does not simply exist. People in organizations produce it, and they do so in regular, patternlike ways.

In addition, once they produce it, they integrate their knowledge into the distributed knowledge base of the organization—also in regular, patternlike ways. To help better understand the combination

of such patternlike knowledge production and integration processes, *the new KM has given us the Knowledge Life Cycle, or KLC—a descriptive model of knowledge processing in human social systems that is unique in the field* (see Figure 11.3 above).

The origin of the KLC (see Chapter 2 and just above) can be traced to individual and organizational learning theories and also to theories about how living systems learn on an emergent basis as expressed in complex adaptive systems (CAS) theory. In addition, we see aspects of the KLC that are deeply rooted in psychology and cognitive science (personal knowledge), social psychology (sense making), and epistemology. When we put all of this

together, we see the manner in which decision making and action are tied to sense making and the use of existing knowledge, as well as sense making and the production and integration of *new* knowledge. These are processes, not outcomes, and they can be strengthened, reinforced, and enhanced. As one of us (McElroy 1999) once remarked (to himself), "It's not knowledge management, stupid, it's knowledge *process* management."

The difference between information management (IM) and knowledge management (KM) was analyzed in Chapter 3. There we concluded that the essential difference between the two is in the area of Knowledge Claim Evaluation (KCE). KCE is a primary target for KM, but it is not logically required by the definition of IM. Moreover, the outcome of KCE is artifact-based knowledge that is distinguished from information by the tests and evaluations performed in KCE. Thus, *the management of KCE lies at the heart of the new KM because it (KCE) is the immediate source of artifactbased knowledge and its outcomes are the basis for the distinction between information and knowledge. Without management of the KCE there is only IM, not KM! With management of KCE, KM presents us with its own autonomous foundation, both as a process and as a discipline.*

In Chapter 3 we also raised the question of how KM may be specified in the form of a framework that offers a middle ground between definition and measurement and that provides a platform on which measurement may be based. We answered this question by providing a framework based on Mintzberg's (1973) classification of executive activities, the distinctions among levels of KM activity, KLC targets, social/technological, and policy/program dimensions. This framework provides the most extensive foundation for classifying KM interventions yet proposed

Supply- and Demand-Side Knowledge Processing

Once the KLC starts to come into view, we can make a further distinction (see Chapters 4 and 5) on how its presence in organizations relates to what we do in commerce—that is, its role and impact on our business processing affairs. When we engage in business processing, we engage in knowledge use. Turning to the Decision Execution Cycle once again (see Figure 11.2), we can see that pre-existing knowledge is used by us as a precursor to action whenever we monitor, evaluate, plan, and make decisions. On occasion, however, as we explained in Chapter 2, our preexisting knowledge fails us. This gives rise to epistemic gaps—gaps in what we know versus what we need to know. Cycles of knowledge processing driven by the learning incentive system then follow in order to close such gaps.

Once inside knowledge processing (the KLC), we can see that one cluster of goal-directed activity is focused on knowledge production (problem solving), while another concentrates on knowledge integration. Here we sometimes find

it useful to refer to these two clusters in a shorthand manner by referring to knowledge production as "demand-side" knowledge processing and knowledge integration as "supply-side" knowledge processing. Why? *Because knowledge production is triggered by the demand for new knowledge in response to epistemic gaps; knowledge integration, in turn, is focused on the supply of knowledge only after it has been produced.* Knowledge integration is all about the propagation, diffusion, and supply of existing (old or new) knowledge across personal and organizational boundaries.

This understanding of demand- and supply-side Knowledge Processing (KP) quickly leads to a similar distinction we can make between demand- and supply-side KM (see Figure 11.4 next page). Demand-side KM consists of KM strategies and interventions aimed at enhancing demand-side knowledge processing, or knowledge production; supply-side KM consists of KM strategies and interventions aimed at enhancing supply-side knowledge processing, or knowledge integration. *In truth, the best KM strategies are ones that cross both boundaries and which deal comprehensively with the whole life cycle, not just Parts of it.* Nonetheless, we can make the demand- and supply-side distinctions, and it is often useful to do so as we consider narrower, more targeted interventions in our work.

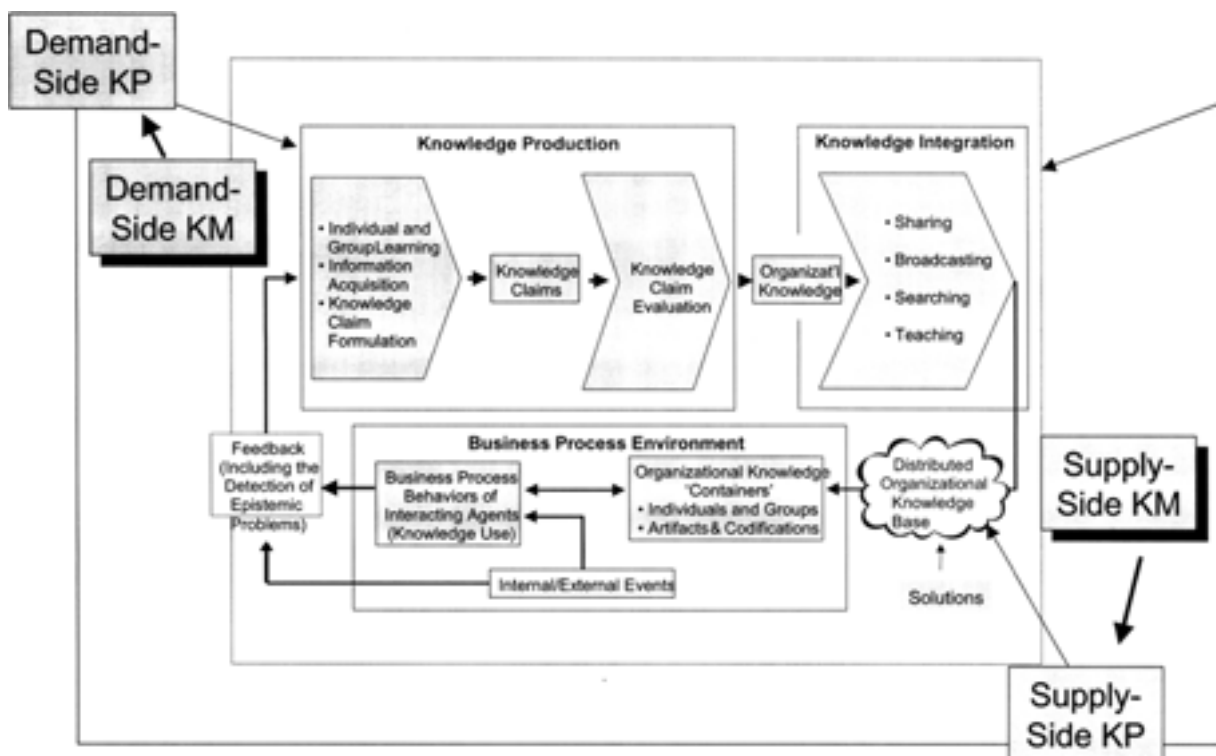


Figure 11.4
Demand- and Supply-Side Knowledge Management

Metaclaims and Best Practices

That the new knowledge management is so focused on the importance of knowledge claims and metaclaims is unique in the field. No other brand or style of practice in KM shares this perspective. Still, the idea seems obvious to us. Why shouldn't we think of knowledge that we express in, say, objective linguistic form as being anything other than a claim? To proclaim something as true is merely to claim it as such (PRO-CLAIM). And as a claim, an assertion is of course not necessarily true, for which of us can also claim to have direct contact with the truth, direct knowledge, or omniscience of a sort? No, all knowledge claims, as well as metaclaims, are fallible.

Adhering to fallibilism, we also agree with Popper that we can make choices between statements or claims that we believe to be true and others that we don't. *And we can do this through the use of deductive logic, which exposes inconsistency (see Chapter 5) and forces us to choose between premises and conclusions of arguments and therefore to falsify some of our statements and grow our knowledge.*

The practical implications for KM and for people operating in business processing modes of the idea that we can falsify our knowledge claims and to select among them from the survivors are enormous. Consider the case of "best practices," for example, as we did in Chapter 7. What is a codified "best practice" if not merely a knowledge claim? It is a claim of truth about what patterns of business processing behavior will fetch the best or most desirable results. But what are its premises? What makes its advocates think so? And why should we, as would-be users of such claims, accept them as true or as close to the truth? Where are the claims that lie behind the claims-where are the metaclaims?

The inescapable fact of such great importance to KM and to the new KM, in particular, is that *people naturally rely on metaclaims for evaluating knowledge claims as a basis for choosing to accept some knowledge claims over others (or not), even if the systems that support such decisions do not explicitly support metaclaim production, integration, or use.* In other words, the processes we use to evaluate knowledge claims, such as "best practices" claims, are always broader in scope than the systems we have historically relied on to support us have been.

A conventional IT-based best practices system will always report claims, but rarely will it provide us with the evaluative metaclaims that lie behind them. Nonetheless, we still subject such knowledge claims to our own search for the metaclaims behind them, and then we always subject all of that (the claims and their metaclaims) to our tests of consistency, simplicity, projectibility, and other evaluation criteria or perspectives. Rarely do individuals simply accept a

claimed best practice and put it into use without at least asking themselves, does this make sense?

Metaclaims that evaluate knowledge claims, like the black keys on a piano, are the context that matters most in knowledge processing, for how can we safely conclude that "middle C" is where we think it is without them? *I know where it is because of where it sits relative to the black keys behind it, and it always sits thus so.* In knowledge management, this changes everything. *There is no KM in the absence of metaclaims that evaluate, only IM (Information Management). Nor is there any knowledge processing in the absence of metaclaims.* The truth is that people in organizations—indeed, all of us—rely heavily on the content of metaclaims from one minute to the next as we constantly try to sort between claims that we ought to accept and the ones we shouldn't. *Knowledge management systems, then, conceived with the idea of enhancing knowledge processing, must make and support the distinction between claims and metaclaims in order to be useful and complete, and this means that management of the knowledge claim evaluation subprocess is a central fulcrum of KM* (see Figure 11.5).

Knowledge Claim Evaluation

In Chapter 5, we presented a framework for description of knowledge claim evaluation, then a normative theory, called *the theory of fair comparison*, providing a high level conceptual outline of how to do KCE, and last an outline of the requirements and structure of a projected KCE software application. *We believe that this analysis is significant for KM, since very few in the field are writing about knowledge claim evaluation, much less offering descriptive frameworks, normative models, and techniques for implementing this critical activity in knowledge production.*

The idea of "fair comparison" of competing knowledge claims is fundamental to our perspective. We contrast "biased" knowledge claim evaluation with knowledge claim evaluation through fair comparison and assume further that KCE is more effective, in the sense that it fulfills certain success criteria, when it is characterized by fair comparison and less effective when it is characterized by bias. Thus, we believe that KM-induced changes in knowledge processing rules and criteria that increase the degree of *fair comparison* also increase KCE effectiveness, and changes that increase the degree of *bias* decrease its effectiveness.

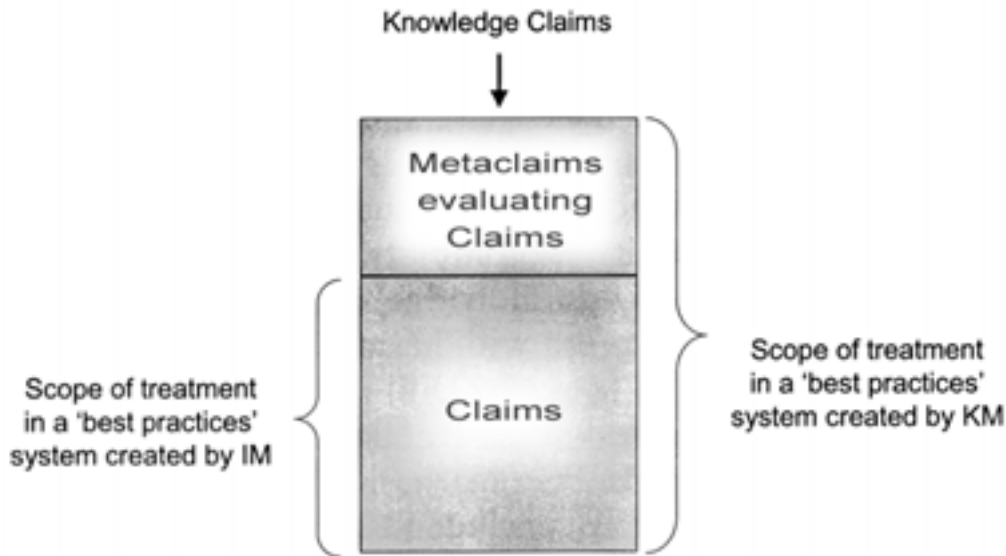


Figure 11.5
Claims versus Metaclaims

Normatively, of course, one should seek to increase KCE effectiveness and therefore increase the degree of fair comparison. We believe this can be done at the level of knowledge processing by:

- First, fulfilling background requirements (the necessary conditions) for fair comparison among the members of a set of competing knowledge claims; and
- Second, implementing comparisons among the members of this *fair comparison set*, based on a number of criteria that allow us to choose among the knowledge claims of the set based on how its members perform on various tests.

The theory of fair comparison specifies *equal specification of members of the comparison set, continuity, commensurability, and completeness of the comparison set* as four necessary conditions of fair comparison. It also names *logical consistency, empirical fit, projectibility, systematic fruitfulness, heuristic quality, systematic coherence, simplicity, and pragmatic priority* as criteria for evaluation of competing alternatives once a fair comparison set is constituted. In Chapter 5, we discussed what we mean by each of the above criteria, and also pointed out that in KCE our procedures for combining criteria can range from the very informal to the highly formal.

Informality in combining criteria is what we normally do. That is, when we have a set of factors to be considered in choosing among a set of alternatives in KCE, we most frequently vet the alternatives with others, and may even subject them to a kind of free-for-all critical process and/or weigh them using intuition

and common sense, and then make our decision about which alternatives are false, which ones we are unsure about, and which are true (or at least most "truthlike"). The process may involve considerable critical interaction with others and often may be collaborative, since many perspectives are better than one in appreciating the importance of the various factors in a decision.

Much of the time an informal process of vetting and weighing is also the most appropriate way of combining criteria. It is so because there may be no time for a more formal and systematic approach. Or because the resources may not be available to implement one. Or because what is at stake in the KCE decision may not be important enough to justify one. Or because we need informality to surface criticisms, creativity, and new ideas in KCE. So whether we should, once fair comparison requirements are fulfilled, implement a formal and systematic approach to multicriterion decision making, or an intuitive approach or something in between, depends upon available resources, time, the need for new ideas, and the cost involved, compared to what is at stake in avoiding error. If resources, time, available formal frameworks, and cost are not "right," the appropriate decision method to use in KCE may well be an informal one.

Knowledge claim networks, including metaclaims, have descriptive and valuational aspects to them. They are networks with both descriptive and value interpretations (Firestone 2001, 2003, Chap. 4)

And they may be compared in terms of the priority values across networks of benefits resulting from actions as specified by each knowledge claim network (or theory or model). The criterion attribute of *pragmatic priority* listed above also encompasses relevance. Thus, the greater the benefit specified in a knowledge claim network, the more relevant is the network from the pragmatic standpoint of the consequences of actions in closing gaps between goal states and actual states.

When, during KCE, knowledge claim networks are compared according to their pragmatic priority, we are not engaged in a comparison of epistemic values, but rather one of the estimated costs and benefits specified by each network in the comparison set. In committing to the rejection of knowledge claims as false, and relying on surviving knowledge claims in actions, the risks we take are a combination of (a) the likelihood that our evaluation rejecting particular knowledge claim networks is in error, and (b) our assessments of the benefit/cost consequences of such errors-and that, as a result, we might suffer the consequences predicted by the true knowledge claim network we have rejected. Thus, pragmatic priority requires that epistemic criteria be weighted by the risk of error in developing a comparative evaluation of knowledge claims and knowledge claim networks. This criterion does not involve wishful thinking, in the sense that we will value most highly those knowledge claims that predict the greatest benefits, but rather modest pessimism in that epistemic

values are reduced based on the risk of error involved in not rejecting the surviving knowledge claim networks, and in rejecting their alternatives.

We expect that the set of direct comparative evaluation criteria we offered in the normative model may well be incomplete and that some of the model's components may be incorrect. So applications of the model should be freewheeling and adaptive, and modifications of it should be undertaken freely. Moreover, the issues of the form of the composite models and the weighting used in them are left open by us, even though we provide a couple of examples of combination approaches in Chapter 5. We believe that these aspects of the normative KCE model will vary across enterprises and that until we have much more experience in applying the model it would be premature to suggest generalization of these aspects of it.

In short, we believe that the criteria, the form of KCE composite evaluation models, and the weights used in such models must all be open to innovation themselves, especially in an environment where we know little about the details of formal methods for evaluating KCE processes.

Thus, we present the normative framework of Chapter 5 as a working theory of KCE, as a first effort in a direction that we need to take in knowledge management. We hope and believe that though this may be the first attempt at a normative model in the KM literature, it will be far from the last, and we look forward to having others join us quickly on this important road not previously traveled in knowledge management.

The Centrality of the Knowledge Life Cycle

What is perhaps most surprising about the old, first-generation KM is its failure to specify the nature or behavior of its target in any sort of systematic or rigorous way. Worse yet may be its failure to acknowledge that a target of its interventions exists. What do we mean by this?

In the new KM, when we make a KM strategy or engage in an intervention of some kind, we are always seeking to enhance the performance of knowledge processing in human social systems. The target of our interventions, that is, is a social process that we are trying to improve. Moreover, the social process of interest to us has a particular shape and content. It has internal elements, subprocesses that we understand in advance of our efforts and which we can target more narrowly if we choose (see Figure 11.6). This social process that we speak of is, of course, the Knowledge Life Cycle, or KLC.

In the new KM, the KLC is of central importance. It is as important to us as practitioners of KM as knowledge of the patient is to a physician. And just as patients and their illnesses figure prominently in the practice of medicine, so do the KLC and its illnesses figure prominently in the practice of the new KM. It is our job, that is, to ensure the health of the KLC and to enhance its

performance whenever we can. And just as a patient in a hospital has a mind and a body for physicians to work with, so does the KLC have substance and regularity to its form. We can see it, touch it, measure it, anticipate its behaviors, alter its behaviors, change the conditions in which it operates, tinker with its internals, and experience its outcomes. In a very real sense, the KLC is a living system, and it falls to the new KM to ensure its well-being.

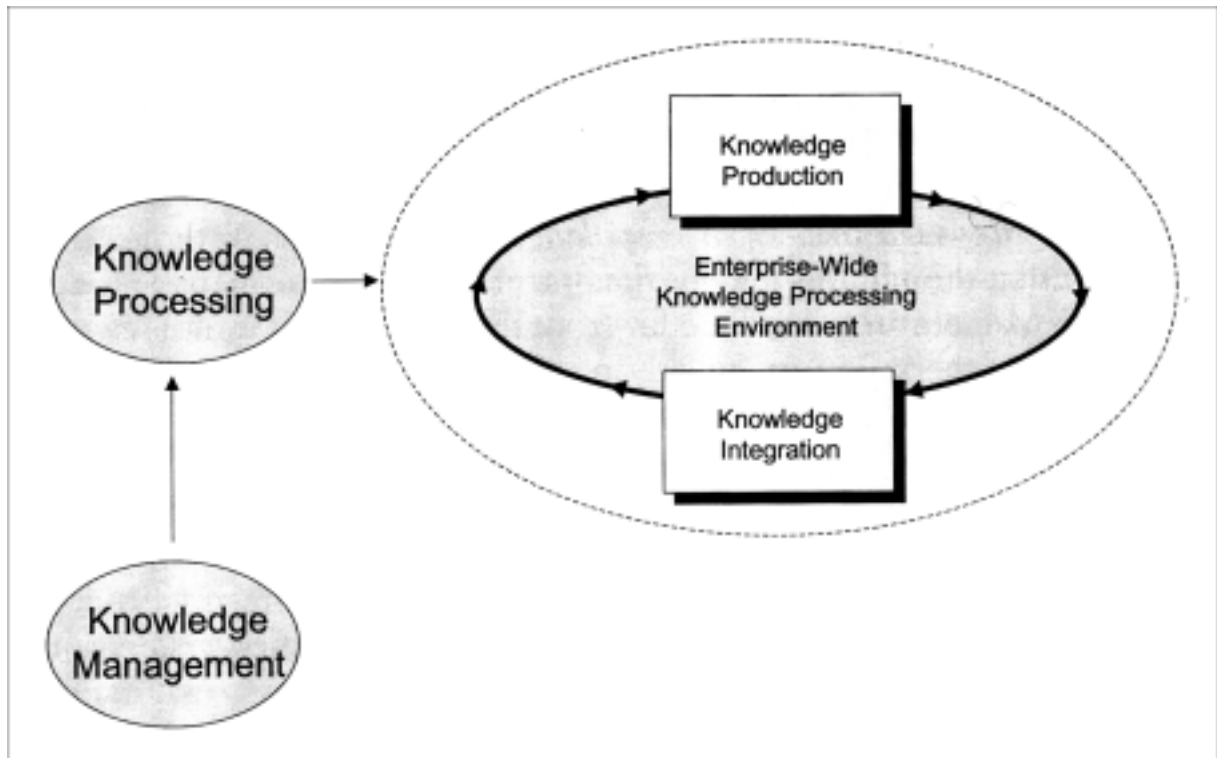


Figure 11.6
Knowledge Management versus Knowledge Processing

What, then, could the equivalent of the "patient" possibly be in the practice of first-generation KM? For all of the obsessive talk of enhancing knowledge sharing that occurs on that side of the fence, we can only guess that it must be the act of offering someone else one's knowledge. Or is it the act of accepting or retrieving such knowledge once offered? In any case, in first-generation KM, there is no distinction between knowledge management and knowledge processing. There is no social process that comprises knowledge processing. There is no model or framework of knowledge processing that can serve as a target of KM interventions or strategies of any kind. There seems only to be life in organizations that is periodically punctuated by sudden needs for information (not knowledge, mind you), backed up by constant browbeating (or stories, or shallow incentives) from management to get people to share it.

Indeed, in the old KM, the *raison d'être* for KM strategies and interventions seems singularly aimed at knowledge sharing and use events that, from the

perspective of the new KM, are isolated transactions in a much bigger knowledge processing system.

And while the new KM sees knowledge processing as the heart of the adaptive engine that makes organizations tick and helps them to survive, the old KM seems inexorably fixated on only their transactional needs. Ensuring that individual acts of sharing and information retrieval occur as best they can is the sole priority of the old KM. Never mind that the firm may be sinking fast, or that managers are looting the treasury, or that the broader adaptive performance of the company is trending downward. There is no higher purpose in the old KM, no KLC, no knowledge processing system, no complex adaptive system, no decision execution cycle-none of that. There are only people who have information and others who need it.

If in the new KM, we can say that the Knowledge Life Cycle lies at the heart of its practice, in the old KM it must be the Information Delivery Cycle (IDC). The IDC has two parts to it: (1) information sharing and (2) information retrieval. And if this is true, then what is the difference between information management and knowledge management in the old KM when all is said and done? The answer is unclear. But in the new KM, the makeup and purpose of the patient is clear. The organization is a complex adaptive system, and the knowledge life cycle is its adaptive engine. It has a "physiology" with recognizable regularity to it that we can clearly observe, understand, and have impact on. Our purpose, then, is equally clear-to enhance the health and performance of the KLC!

In Chapter 6 we illustrated the centrality of the KLC by illustrating various areas of its application. There we discussed how the KLC is applied to KM strategy, KM and knowledge audits, modeling, predicting, forecasting, simulating, impact analysis and evaluation, metrics segmentation, sustainable innovation, methodology, IT requirements, intellectual capital, education and training, the open enterprise, and new value propositions for KM.

There are many other areas of interest to practitioners of KM where the KLC has a valuable role to play. *The scope of its influence on the practice of KM is extensive, and our discussion in Chapter 6 focused on only a few of them. Nevertheless, the breadth and significance of the applications discussed illustrate the centrality of the KLC to KM and also its utility from the point of view of practice. It also illustrates that the KLC changes everything for KM. Its implications penetrate every nook and cranny of KM as a discipline and every area of KM practice.*

Again, the importance of the role that a conceptual framework or model of knowledge processing plays in the practice of KM cannot be overstated. It is absolutely critical that we begin the practice of KM with a framework or model describing in broad outline how knowledge is produced and integrated in human social systems. From there we can begin to formulate practice, but always with reference to our theory. We must have grounding. All bodies of practice in KM should be held to this standard. What are their interventions?

What methods do they use? How are their methods and interventions related to knowledge processing? Indeed, what is their theory of knowledge processing? Do they even have one? If not, how do they account for knowledge production and integration, and how does their practice relate to their theory? If not, why should we accept their value claims?

As we sometimes say, the only thing worse than all theory and no practice is all practice and no theory. *In truth we need both—practice grounded in theory.* In the new KM, the KLC is the most important focus of our theory. Many of our practices, then, are grounded in and upon it, and from the point of view offered to us by the KLC, the mission of the new knowledge management—to enhance knowledge processing—is clear!

KM and Strategy

Another of the most common mistakes made in the practice of the old KM is what we call the "strategy exception error." The strategy exception error mostly shows up in the composition of first-generation KM methodologies. *The error is committed by turning to existing business strategy as the starting point for KM strategy.* The logic behind the error is that (1) the purpose of KM is to help fulfill strategy (business strategy), and (2) that all KM interventions must therefore be planned by first making reference to the content of current business strategy. The error continues, then, by making yet another one—that there is no difference between information and knowledge (no claims and metaclaims) and that the proper role of KM is to improve information delivery in the service of strategy.

KM methodologies that commit this error usually go on to specify steps that run something like this:

- Step 1: Identify current business strategy.
- Step 2: Determine information resources required to successfully carry out current strategy.
- Step 3: Perform IT and other organizational projects required to make information resources easily accessible and supportive of business processing.

The result of course is a transactional one. In other words, at best KM strategies carried out according to this methodology wind up improving fleeting acts of information sharing and retrieval, but not necessarily anything else. Here again, this has the effect of reducing KM to nothing more than information management trotted out in today's more fashionable clothes, even as it breeds ill will in the business community for what appears to be a grand deception, or at least an irresponsible form of self-delusion. We're here, of course, to point out that the old KM emperor, in fact, has no clothes.

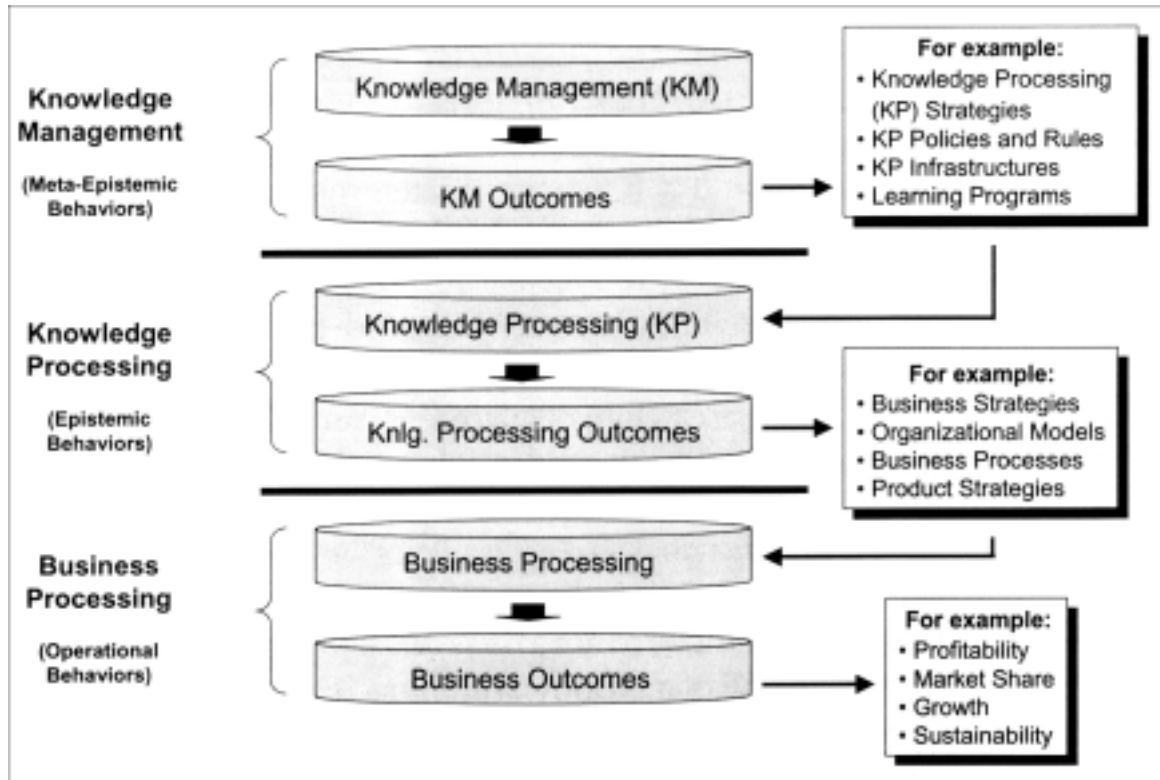


Figure 11.7
The New Knowledge Management Reference Model

What, then, is the relationship of strategy to KM? Quite simply, strategy is an outcome of knowledge processing. So in order to answer the question, one must first be equipped with an understanding of the difference between business processing, knowledge processing, and knowledge management (see Figure 11.7). Business processing is the domain of business strategy and related behaviors. *Knowledge processing, however, is where business strategies are hatched. They (strategies) are the manifest resolutions of epistemic gaps between what managers in organizations need to know about strategy and what they, in fact, do know. Thus, a business strategy is a network of knowledge claims produced as a consequence of knowledge processing.*

Knowledge management, in turn, IS not subject to strategy, nor is it supposed to help fulfill it. That role falls to information management and to operational business processes, not KM. KM's role transcends strategy. KM's role is to enhance the quality and performance of knowledge processing behaviors, not business processing ones. Indeed, the quality of a business strategy depends heavily on the quality of knowledge processing, its source, and thus indirectly, as well, on the quality of KM. But strategies, like all knowledge claims, are fallible; they come and go. KM, though, like good financial man-

agement practices, is permanent and transcends strategy. It's above and apart from strategy, not within or below it.

This perspective on the proper role and positioning of KM in a firm raises another important issue for which the new KM has an emphatic answer. *The question is where KM belongs in the functional configuration of a company. The answer, we believe, is that it should be reporting directly to the board of directors and to no one else. Why? Because maintaining the independence and integrity of knowledge processing in a firm is a fiduciary matter, just as the quality of financial management is. Oversight for both, therefore, is best placed at the level of the board.*

Anyone who has any doubt about this need not think for long about what's happened in the United States over the past year or so. When we look at the crises that occurred at such firms as Enron, Worldcom, ImClone, Tyco, and others, what we see are bad strategies, or parts of them, in use. What we can ask ourselves, then, is what must the knowledge processing systems be like in such firms that knowledge claims *of such poor quality manage to survive their internal tests and evaluations?* Who was watching the knowledge processing store there? Was anybody watching it?

Perhaps it can be said that one consequence of the information or knowledge age that few of us fully appreciate is that with the ascendance of knowledge as the new prized asset in corporations has come the need to exercise oversight on its production. *The time has come for boards in businesses (especially publicly traded ones) around the world to formally embrace oversight for knowledge processing as a fiduciary duty, and to recognize the fact that the knowledge we practice in organizations is . is the knowledge we produce. Boards must rise to this challenge, and the new KM can help show them the way.*

KM and Culture

What is culture, and what is its relationship to knowledge and knowledge management? "Cultural" barriers are often held responsible for failures to share and transfer knowledge in organizations. It is frequently said that knowledge management must undertake the difficult task of changing an organization's culture to achieve the knowledge sharing and transfer necessary to realize the full value of the organization's knowledge resources. In Chapter 9, we showed that this viewpoint is incorrect, and that the widespread belief in it is probably due to the use of the term to describe many different factors that are social, political, or psychological, rather than cultural in character. Our analysis produced the following definitions of culture.

The subjective culture of a group or organizational agent is the agent's characteristic set of emergent high-level predispositions to perceive its environment. It includes group or organizational level value orientations and high-level attitudes and the relations among them. It is a configuration of

global attributes that emerges from group interactions-that is, from the organization and pattern of transactions among the agents within a group. *The objective culture of a group or organizational agent is the configuration of value orientations and high-level attitudes expressed I . in the agent's characteristic stock of emergent problems, models, theories, artistic creations, language, programs, stories, etc., reflected in its documents, books, art galleries, information systems, dictionaries, and other containers.* It is a configuration of global attributes expressing the *content* of its information, knowledge, art, and music, apart from both the predispositions the group or its agents may have toward this content, and the material form of the artifact expressing the content. The objective culture of an organization is an aspect of the social ecology of its group agents, consisting of cumulated effects of previous group interactions. As such, the perception of it by individual and group agents (part of their subjective, culture or psychology, depending on the type of agent) influences~ their behavior.

As we argued in Chapters 2 through 4 and just above, we can distinguish KM processes, knowledge processes, and business processes. And knowledge processes may be viewed in terms of the KLC framework. KLC processes produce knowledge that is used in the other~ business processes of the enterprise. And these, in turn, produce business outcomes. Figure 11.8 illustrates this chain of influences.

Moreover, KM processes, knowledge processes, and business~, processes are performed by decision-making, behaving agents. As we have seen, agents, if they are groups, have an internal culture, both subjective and objective. At the same time, the objective cultural component of social ecology also impacts agent decisions. Finally, knowledge and KM processes are affected by culture through the influence it has on behavior constituting these processes. In turn, these processes are keys to producing new knowledge and consequently changes in objective and subjective culture. The interplay between KM processes and culture is bidirectional.

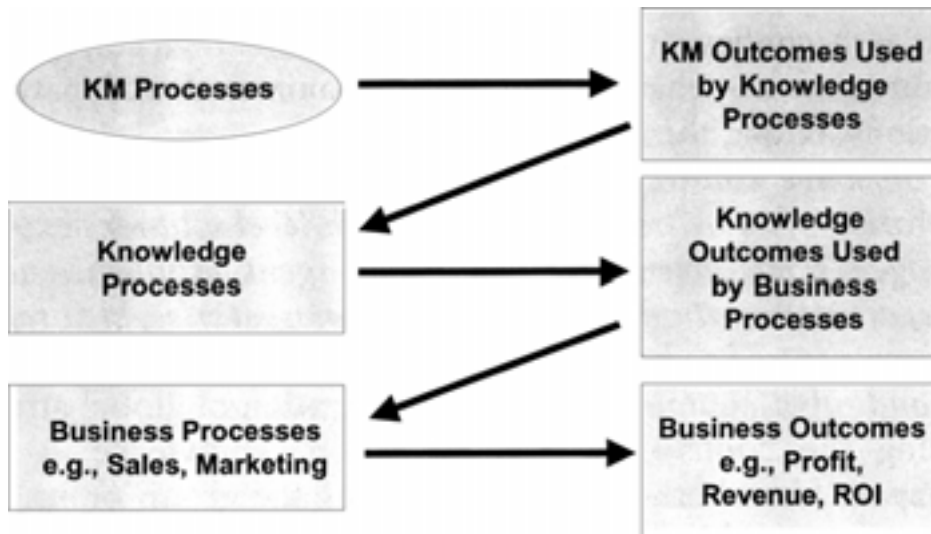


Figure 11.8

From Knowledge Management Processes to Business Outcomes

So culture is pervasive in KM, knowledge processing, and knowledge outcomes. It IS part of their context, and it is also, in the long run, produced by them. But many other factors (social ecology, situational factors, transactional inputs; see Figure 11.9) also contribute to the complex interactions associated with knowledge-related processes and outcomes. Thus, culture is only a small part of what there is to KM, or any other business process, and therefore there remain substantial problems in measuring and analyzing its precise impact on KM, or KM's impact on culture. Culture is not so much an answer to difficulties in KM as it is an issue and a problem in I . itself And prescriptions that suggest that we must "change the culture" to perform effective KM and to enhance knowledge processing are not solutions to problems, but only prescriptions for movement down the wrong track toward a single-factor explanation of knowledge management and knowledge processing.

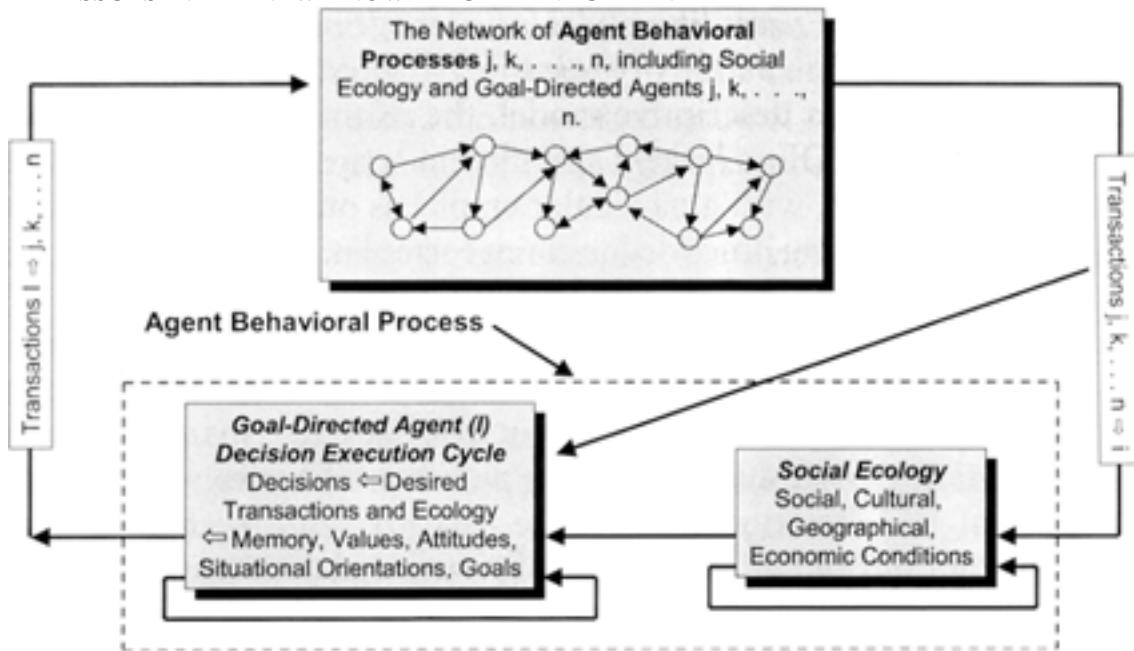


Figure 11.9

The Flow of Behavior Among Agents and the Role of Culture

The Open Enterprise

This is perhaps the most exciting idea to come out of the new KM.

In a phrase, *the Open Enterprise (OE) is a normative model for knowledge processing and knowledge management designed to achieve sustainable innovation and transparency in management. It is the antidote to Enron-like styles of management even as it also enhances innovation and the organizational capacity to adapt.*

Unlike the KLC, which is a descriptive model, the open enterprise is a prescriptive one. The OE is largely specified in terms of the KLC's subprocess dimensions, with a particular emphasis on problem detection in business processing, knowledge claim formulation, and knowledge claim evaluation.

What makes the OE so exciting to us is that it is the first comprehensive, normative vision of knowledge processing to appear on the KM landscape. As such, it can equip knowledge managers with a specific target environment, a to-be picture of what to shoot for in the area of KM strategy and practice. *It therefore provides us with a description of how knowledge processing ought to be in a firm, in terms specific enough that we can translate them into KM strategy and action.* And unlike the narrow, transactional bandwidth of the old KM, according to which the to-be environment is confined to individual acts of information sharing and retrieval, the OE is system- and process-wide in scope. It offers a vision of what highperformance knowledge processing should (or can) look like in all of its dimensions and at all level s-enterprise -wide.

The Open Enterprise has several important value propositions not found elsewhere on the KM landscape. First, it enhances an organization's capacity to detect epistemic problems by effectively engaging the whole firm in the process, not just the designated elites in management who can't possibly have the full range of view required to detect problems on a distributed basis. Second, it implements its policy of broad inclusiveness by enfranchising all of the organization's members in the knowledge production side of the KLC, especially in terms of knowledge claim formulation and evaluation. Thus, the full human resource complement in a firm, as opposed to only a fraction of it, becomes engaged in knowledge processing.

Also different in the OE are the rules that relate to knowledge processing. No longer is information or knowledge held close by management hierarchies. Instead, knowledge processing itself is more distributed as is the previous knowledge that informs it. In the OE, the separation of business processing and its management from knowledge processing and its management are made explicit. The monopoly on knowledge processing control once held by business processing managers is repealed. Knowledge processing becomes open to all stakeholders in a firm, not just its managers. This, however, does not mean that management is "undermined." To the contrary, the power to control and direct resources remains with management on the business processing side of the equation.

Considering the issue of terrorism, to be discussed in more detail below, we believe that the specific contribution the new KM can make to the quest for security in the United States and elsewhere is to show how knowledge processing can be improved by bringing powerful new normative (i.e., prescriptive) models, such as the OE, to the table. Only when knowledge processing systems are fully open in terms of the kinds of structural and operational attribute values present in the OE pattern will intelligence be capable of operating fully at its peak. Had the U.S. intelligence system been operating in accordance with our vision of the OE, it may well have been able to "connect the dots" long before the attacks of 9/11/01 took place. We hope to substantiate this claim in more detail as the full scope of the OE is developed in the months ahead.

Intellectual Capital

Underlying all of our comments in Chapter 10 regarding the unsuitability of conventional accounting perspectives in the treatment of intangible values are four key points:

1. The most important source of intangible economic value in a firm is external to itself-not found at all within the firm, and therefore not directly manageable by it. The origin of such value is in the minds of current and prospective stockholders, whose valuations of a firm are determined within the dynamics of their system, not the enterprise's. Their values are then projected onto the enterprise, which in turn receives it as a kind of gift or inheritance-an inheritance, however, that is always subject to adjustment, and even return.
2. Apart from being determined by agents outside of the enterprise, reflecting the value of intangibles using conventional accounting tools as though they were internally managed suffers from another problem, as well. Balance sheets are static models. They are predicated on the assumption that it is possible to start with a market value that can then be broken down into its additive elements. *Knowledge and the processes that produce and integrate it, however, are anything but additive. What we need in order to account for them all are dynamic models-nonlinear ones-not static ones.* These, in turn, should come into play in conjunction with the theories of emergence they support, the composition of which will suggest reporting models and tools of their own. This is the path to solving the so-called IC problem. Ramming new realities into old frameworks simply will not work.
3. All of our warnings concerning the inability to manage emergent systems and outcomes have been precautionary, but not absolute. *While it's true that processes and outcomes in complex adaptive systems are emergent beyond the absolute control of managers, it is possible to have impacts on them,, nevertheless-even intended impacts that give rise to corresponding (hoped for) outcomes.* This, however, does not entail conventional management. Rather, management in such environments must begin with recognition of organizations as social complex adaptive systems, subject to the mostly unpredictable effects of system dynamics and emergence. Armed with this insight, it is possible to build models of organizations and their environments that are nonlinear in form and to use such models for predicting the potential outcomes of alternative management strategies, including the effects that management decisions of one kind or another might have on the externally produced value judgments assigned by markets to the enterprise. In this way, managers can have impact on the value of their organizations' organizational and market intangibles, which, so far as we're concerned, makes such values *manageable*.
4. That said, we are only prepared at this point to support the reflection of intangible values in the lump-sum fashion we have proposed. It is simply too soon to suggest any other, more granular treatment of intangibles on

the balance sheet. *What's required, instead, is a period of trial and error, or experimentation, . . . on, in which competing theories of how intangible values are produced can be tested and evaluated.* As a consequence of that process, we can conceive of a point in time when the momentum behind one or more models of how intangible values are produced matures, and informs us of their management, measurement, and reporting implications. Then and only then will we have a solid basis for suggesting how best to report on the value of intangibles to stockholders. Until then, it's premature to attempt such a report.

On the basis of these remarks, we believe that the accounting profession should undertake a concerted effort to bring additional, nonaccounting disciplines into the process of trying to discover the nature of intangible value and the role that it plays in corporate valuations. These would include complexity scientists, social scientists, and knowledge managers of the second-generation (new KM) type. Until we do, and until we find ourselves testing and evaluating the use of nonlinear models, not linear ones, we should expect, as Kuhn (1970) put it, to encounter nothing but problems that seem to indefatigably resist "the reiterated onslaught of the ablest members of the group [accounting] within whose competence it [the IC problem] falls."

Information Technology and the New KM

No summary of the landscape of KM (new, old, or otherwise) would be complete without making some remarks about the IT implications of related practice. In the case of the new KM, the story, we're afraid (see Firestone 2003, Chaps. 10-19), is a short one. *At present, there are no software applications that we know of that explicitly address the claim /metaclaim distinction so central to the new KM, or provide explicit support for knowledge claim evaluation.*

In particular, the Enterprise Knowledge Portal (EKP) (Firestone 1999, 2000a, 2003, Chap. 13) is an application that comprehensively supports knowledge processing and knowledge management. *Despite claims to the contrary, such an application does not yet exist. . . . But the last two years have brought progress toward achieving this essential objective of TNKM.* Firestone's (2003) review of various KLC and KM categories of TNKM framework provides a view of the gap between the current state of portal progress and what is needed for an EIP that would support knowledge processing and knowledge management—that is, for an EKP.

Specifically, current EIPs provide support for those subprocesses in knowledge processing and knowledge management that are common to KM and information processing and information management. However, they don't

support individual and group learning, knowledge claim formulation, knowledge claim evaluation, knowledge outcomes, the DOKB, KM knowledge processing, resource allocation, or negotiation well.

The most glaring departure from TNKM requirements IS in the area of knowledge claim evaluation. Here, current portal products provide almost no support and the idea of providing it is apparently not even on the radar screen of any vendor Perhaps that will change. But we are still probably years away from having a real knowledge portal. In Firestone (2003, Chaps. 13, 17, and 19), we outlined "how to get there from here."

Slow progress is perhaps due to the relative obscurity of ideas about the KLC, KCE, and metaclaims in conventional KM circles. We hope, of course, that this book and other writings of ours will help to change this and that progress will be made in recognizing the centrality of the KLC and KCE and the role that claims and metaclaims play in our knowledge processing affairs.

That all said, the landscape here is not entirely barren. Using the KLC as a backdrop, we can easily see existing technologies and applications that support aspects of knowledge processing. E-learning systems, for example, clearly support individual learning on the knowledge production side of things-though not yet very well in a work-oriented problem-solving context-as well as training on the knowledge integration side. Similarly, communities of inquiry that support at least the human interaction, if not the recording and tracking side of error elimination in knowledge processing, can be supported by a wide range of tools aimed at collaboration and group learning. And certainly the information acquisition subprocess is extensively supported by a whole host of applications, some of which have been around for some years.

Of particular interest to us, however, are the new technologies that support the analysis and management of knowledge claims in artifacts. Tools that rely on semantic network analysis, neural network analysis, fuzzy logic, and other technologies aimed at teasing out claims and the metaclaims behind them in e-mail messages, documents, and other electronic files are particularly promising here. But the best in new KM technologies is yet to come, as we discuss in the next section on the future of the new KM.

THE FUTURE OF THE NEW KM

In Popperian epistemology, the development of new knowledge always raises new problems, and therefore a need for more new knowledge. The new KM is no exception. As a new science, much of its content and methodology is still not fully defined, and many of the new questions it raises have not yet been answered. In some cases, these are entirely new questions; in others they are old questions answered by the old KM, for which the new KM has new answers to offer. What we offer here, then, are some brief discussions of new KM

perspectives on some new and old problems in KM. In forthcoming works of ours, we plan to address each of these issues in more detail. For now, simply naming and describing the problems to be worked on will do.

SECI Model

Among the many important implications of the new knowledge management is that the once, and perhaps still, popular SECI model put forth by Ikujiro Nonaka in his 1995 book with Hirotaka Takeuchi, *The Knowledge- Creating Company*, may be, as we indicated earlier, materially incomplete and seriously flawed. While the reformulation of Nonaka's model by us is not yet fully completed, we can say that the original SECI model (Socialization/ Externalization/Combination/Internalization) suffers from two important oversights.

First, the SECI model has many flaws at the level of psychological and cognitive theory. Thus, it neglects to include consideration of implicit knowledge and in the process provides us with an ambiguous rendition of tacit knowledge. As Michael Polanyi, an early source of theory for the nature and meaning of tacit knowledge, pointed out, some tacit knowledge consists of that which one can know but never tell- ". . . we can know more than we can tell. " (Polanyi 1966, 4). Tacit knowledge is, therefore, inexpressible. There is no conversion from it to explicit form, thus Nonaka's notion of externalization (acts of converting tacit knowledge to explicit knowledge) is misleading. Implicit knowledge, however, can be converted to explicit form. But Nonaka and Takeuchi made no provisions for implicit forms of knowledge, even though Polanyi mentions it in his work (Polanyi 1958, 286-288).

Further, the SECI model also fails to distinguish between knowledge predispositions and situational orientations. The distinction between tacit and explicit knowledge may either be interpreted as applying to predispositions or to orientations. But if it's applied to predispositions, it has no meaning because there are no predispositions that are explicit. On the other hand, if tacit knowledge is applied to orientations, then it is clear that much of the "tacit knowledge" that people have referenced in examples, such as the ability to ride a bicycle, doesn't fit such a notion of tacit knowledge, because such abilities are predispositional in character.

Second, with the arrival of the new KM comes the distinction between subjective knowledge in minds and objective knowledge in artifacts. This, too, materially expands the range of possible forms from which knowledge can be converted. Instead of just tacit, implicit, and explicit forms of knowledge, we now have all three for *both* subjective and objective knowledge to consider. Thus, we have up to six categories of knowledge to deal with in the new KM knowledge conversion model, not just two as in the case of Nonaka's SECI model. In truth, however, there are only five such categories since we have

already rejected the notion of tacit objective knowledge, or tacit knowledge held in anything other than minds.

All of this leaves us with a dramatically expanded matrix, which takes us from a two-by-two matrix with four cells to a five-by-five matrix with twenty-five cells. The implications of this insight are enormous. *Among other things, it means that all of the KM practices and practitioners out there that have been basing their knowledge conversion efforts on the SECI model may have been engaging in fantasy for the past several years. There is no conversion of tacit knowledge to explicit knowledge; there never has been and never will be.* At best, they've been dealing with implicit knowledge, not tacit knowledge, but the SECI model never allowed for that. Here we see a dramatic example of how theory matters greatly to practice. Faulty theory can lead to faulty practice, which can, in turn, lead to colossal wastes of time.

In sum, then, the landscape of ideas contained in the new KM raises several important questions about the falsity and usefulness of the Nonaka SECI model, thanks to the new KM's identification of implicit knowledge, objective knowledge, and subjective knowledge as forms of knowledge that the SECI model simply overlooks. In a future work of ours, we plan to unveil a fully reformulated *knowledge conversion* model that corrects these errors-which, of course, will be- nothing more than a new knowledge claim complete with its own set of new questions and epistemic problems.

The EKP

A vital task for the new KM is development of the Enterprise Knowledge Portal (EKP), because it is that application that provides comprehensive support for both the KLC and for KM. In addition, the EKP is vital for the open enterprise as well. In Firestone (2003, Chap. 13), one of us outlined the steps that should be taken to advance from present EIP platforms to the EKP. *The Enterprise Knowledge Portal is an application on the verge of development. The technology it requires is in existence now. The cost of its development IS low as software applications go, since its implementation 1 . is largely a matter of systems integration, with the exception of its Intelligent Agent (IA) component that exceeds current IA capabilities. On the other hand, the benefits associated with the EKP are great. They are nothing less than realization of the promise of the Enterprise Information Portal (EIP) to achieve increased ROI, competitive advantage, increased effectiveness, and accelerated innovation.*

EIPs are risky because (neglecting data quality applications which involve relatively superficial quality issues) they fail to evaluate the information they produce and deliver for quality and validity. Nothing, including EKPS, can ensure certainty about information, models, or knowledge claims. But EKP applications incorporate a systematic approach to knowledge claim testing and evaluation that eliminates errors and produces quality assured information. In

the category of portal technology they, not EIPS, are the best we can do. They, not EIPs, are the future of portal technology.

Framework for Analysis of KM Software

The EKP concept provides the basis for a comprehensive framework for evaluating knowledge processing and KM software. Currently there is no such framework. Vendors offer claims about the support they provide for KM, but there is no benchmarking framework against which to evaluate such claims. A future task for the new KM development is to use the EKP concept to produce such a framework.

Role of Credit Assignment Systems in KM Software

The New KM (TNKM) suggests that the primary motivation for using KM software will be the desire to solve problems. However, a long-recognized problem in the old KM has been that of "incenting" people to participate in knowledge sharing initiatives and specifically in software designed to support knowledge sharing. Thus, a natural question is whether the same problem will exist in TNKM. That is, what kinds of incentives will be needed to encourage knowledge worker participation in TNKM software applications? Will credit assignment systems be necessary? Or will it be sufficient to simply involve people in KLCs by giving them free rein to solve their own problems?

TNKM Metrics

In Chapter 6, we discussed the application of the KLC model to metrics development and showed how it could be used as a guide to metrics segmentation. Of course, the KM approach developed in Chapter 3 can also be used to supplement the KLC. Preliminary work has been done in this area by Firestone (2000) and Firestone and McElroy (2002). However, previous work has only initiated the process of KM metrics development. This will be a primary and ongoing activity of TNKM.

TNKM and Terrorism

After the 9/11/01 attacks, much of the criticism levied against U.S. intelligence agencies revolved around their failure to "connect the dots." Critics argued that internal warnings and red flags should have been noticed and that sufficient evidence existed prior to the attacks to suggest that danger was imminent.

From our perspective, the failure of U.S. intelligence agencies to potentially anticipate the 9/11 attacks was a failure of knowledge processing. Using the KLC as an analytical tool, we could say that there was a breakdown in

knowledge claim formulation on a grand scale, but also that on a micro scale certain knowledge claims made by specific agents in the system were too quickly discarded or overlooked. In that case, it was knowledge claim evaluation that failed us, because authoritarian and bureaucratic rather than epistemic and appropriate value criteria were used to falsify promising knowledge claims in favor of others that proved false in the end.

Intelligence is an industry that is in the knowledge processing business. It exists for no other reason than to close epistemic gaps in the knowledge we need about threats to our security versus the knowledge we have. *One could argue, therefore, that there is no more important, more urgent need for the new KM than in the intelligence business, for the KLC provides it, and us, with a roadmap of how knowledge processing happens, and therefore a framework for aiming strategies, tools, and interventions at its improvement.*

Further, it may also be the case that the most urgent software application for an interagency intelligence system is an EKP. Indeed, if the open enterprise can be described as an enterprise-wide knowledge processing system with maximum inclusiveness, then the EKP can be seen as its indispensable IT infrastructure. From the perspective of the EKP, every stakeholder in a firm is a full participant in knowledge processing. If a memorandum written by some far-flung agent in Phoenix, Arizona, suggests that terrorists are going to fly fully loaded passenger jets into the World Trade Center towers in New York City, the EKP picks up on that, connects the dots, and broadcasts its knowledge claims—without fail and without additional human intervention—to appropriate stakeholders in the system. Why? Because the OE is a knowledge-claim-centric construct, not a management-centric one. A credible claim is a credible claim, regardless of who develops it or what their rank or status is in the firm. The knowledge processing ethic of the OE is blind to such things, and the EKP does nothing but doggedly hunt for dots and explore ways of connecting and evaluating them.

The Open Enterprise, Again

The Open Enterprise (OE) is a vital part of the future of TNKM. In future work of ours, we will be developing the model of the OE comprehensively and in much more detail than was possible in this book. Out of this work will come new and potentially powerful applications of the OE for enhancing national security and reducing terrorism, as the approach the OE prescribes will make it possible for intelligence agencies to dramatically enhance their capacity to detect and solve epistemic problems.

The OE ethic of inclusiveness and aggressive knowledge claim formulation will have equally dramatic effects in less urgent settings such as businesses. This, of course, will be its most common application. In businesses, the OE will, by definition, lead to greater transparency and inclusiveness, since its most

important precept is that knowledge processing must be a politically open process. While there will always be a place for secrecy, privacy, and confidentiality in the OE, a great deal of what today passes for closely held knowledge in organizations will be more widely shared in the OE.

The effects? Greater transparency will lead to fewer cases of management malfeasance as inherently bad ideas get swept aside long before they escalate into practice. In addition, the advent of such tools as free employee presses in which employees can openly review and critique management decisions, related strategies, and assumptions will make it easier for stockholders to witness and understand management's actions and intent, as well as the risks associated with both. And finally, the rate and quality of innovation will improve, as more of the firm's employees and other stakeholders become directly involved in problem detection and knowledge production.

Of additional importance to the OE will be normative models for OE-related KM tools and methods. In other words, if the OE is an attractive end-state model, the next question, and one we will devote much attention to in our future work, is *how do we get there from here?* Much work in the new KM has already occurred here, in that a technique known as the *Policy Synchronization Method* (PSM) has been developed and continues to be refined (McElroy and Cavaleri 2000). PSM is a KM method predicated on the view of knowledge processing as a self-organizing complex adaptive system. It therefore concentrates its focus not so much on the system itself, but on its background conditions instead. *Getting a knowledge processing system to function at its peak capacity, according to the PSM method, is best accomplished by managing its background conditions so that they are synchronized with the problem-solving predispositions of people and groups inside of it (the system). Once such conditions are properly set, sustainable innovation and high-performance knowledge processing flourish.*

The importance and significance of the open enterprise cannot be overstated here. The potential contribution that the OE can make to business and society in general is enormous. Anything that can improve our collective capacity to detect problems, eliminate errors in our knowledge, and choose actions that enhance the sustainability of our course is desperately needed at this time. This is the promise of the OE. It is a normative model for knowledge processing in human social systems designed to enhance our capacity to learn, avoid errors, and adapt.

Communities of Inquiry (CoI)

The old KM is much concerned with Communities of Practice (CoPs) because of their usefulness for knowledge sharing. Practitioners in the CoP area often believe that they are useful for knowledge production too. But TNKM perspectives raise the concern that knowledge production in CoPs is likely to be

characterized by the use of consensus as a validation criterion for knowledge claims. This communitarian form of knowledge production is inconsistent with fallibilism and an orientation toward error elimination. Therefore, TNKM must develop an alternative model to the CoP construct specifying the attributes and characteristics of communities dedicated to the discovery and elimination of errors in knowledge that differ from CoPs in behavior and intent. We call such communities Communities of Inquiry (Col). Thus, Cols are the counterpart to the OE at the group or community level.

Knowledge Management Methodology

KM project methodology (Firestone 2001a) is another important aspect of TNKM we have not been able to cover in this work. Clearly, however, it must be a continuing concern of TNKM. In a future work on TNKM, we will take up three TNKM methodology issues. First, should TNKM methodology be a life cycle methodology or should it be an adaptive, iterative, incremental methodology? Second, what implications does TNKM have for the form of that most common of KM initiatives: the knowledge audit? And third, what is the role of the KLC in KM methodology?

Value Theory in Knowledge Management

In Chapters 5 and 10 in passing, we introduced the idea that valuational knowledge claims enter the Knowledge Claim Evaluation (KCE) process. That raises the more general question of the place of valuational models in TNKM. We can approach this question through the KLC. If we do, it is immediately apparent that problem recognition, as well as every subprocess of the KLC, involves making value judgments. However, the subprocess where the value judgments we make seem most controversial is the KCE process. And the reason for that is the legacy of "value free," "objective" inquiry—an idea that was dominant in scientific philosophy for a good part of the last century, but which is far from dominant now.

In TNKM, our notion of objective inquiry views value theory as playing a vital role in the KCE process. Let's review that role and see what it implies. In developing our normative model for KCE we included a category of criteria called *pragmatic priority*. All the other criteria discussed fall into the category of traditional epistemic criteria for comparatively evaluating *factual* knowledge claims. But pragmatic priority involves taking account of the valuational consequences of rejecting knowledge claims as false and relying on surviving knowledge claims as a basis for action.

The risks we take are a combination of the likelihood that our evaluation rejecting particular knowledge claim networks is in error, and the benefit/cost consequences of such errors. If we are in error in our falsifications, we must

suffer the valuational (cost and benefit) consequences predicted by the true knowledge claim network we have rejected. *To take account of these risks in estimating pragmatic priority, we must formulate valuational knowledge claims that provide a value interpretation of our descriptive knowledge claim network. So, to estimate pragmatic priority we have no choice but to formulate a value theory and to use it in making our estimates and in comparatively evaluating factual knowledge claims.*

Now, taking this one step further, the value theory we use to estimate pragmatic priority is only one of a set of alternative value interpretations that might be applied in the specific KCE context in question. And just as we have constructed a fair comparison set of factual knowledge claims that our KCE process must test and evaluate, we also must construct a fair comparison set of alternative value interpretations for testing and evaluation and seek to eliminate error in that comparison set.

So the implication of the role of valuational knowledge claims in KCE that we have just outlined is that objective inquiry requires not only the formulation of valuational knowledge claims, but their testing and evaluation as well. *Valuational knowledge claims are just as important as factual knowledge claims in our knowledge claim networks, and they are just as approachable as factual claims through applying Popper's method of conjecture and refutation focused on error elimination.*

In the case of factual knowledge claims, the goal that regulates inquiry is the truth. *In the area of value inquiry, the goal is "the legitimate."* *But in both cases, the Popperian tetradic schema for problem solving can still be applied to grow our knowledge.* Because in each area, inquiry starts with a problem, continues with tentative solutions, proceeds to error elimination, and then gives rise to a new problem. And the orientation of TNKM and the open enterprise are both just as relevant to the production and integration of value knowledge as they are to the production and integration of factual knowledge.

In the last several paragraphs we have expanded on our account in Chapter 5 and presented a broader perspective on the role of value theory in KCE. That role has implications for the future of TNKM. First, part of our program must be to refine the normative model presented in Chapter 5 and evolve it toward a KCE methodology. Second, we must clarify in future work how value interpretations and value theories are formulated. One of us (Firestone 2001) has performed previous work on value interpretations. Third, we must develop a parallel normative model for evaluating the legitimacy of valuational knowledge claims, i.e., the value interpretations that transform factual theories into valuational theories. And fourth, we must, in TNKM, begin to formulate value interpretations along with our theories about fact.

The New Knowledge Management and Knowledge Management Standards

We close this chapter and the book itself with some final words about standards for KM. We offer our remarks by first acknowledging that there are currently many efforts around the world now under way to develop standards for KM, though generally these are not coordinated. We believe that all such standard formulation efforts are premature at best and seriously misguided at worst. They have the potential of doing great damage to the continued growth and evolution of the field. How? By locking KM and knowledge processing in the industry into patterns of behavior that may, in fact, be harmful to sustainable innovation, organizational intelligence, and organizational adaptation.

Jan Hoffmeister, one of our colleagues on the board of the Knowledge Management Consortium International, points out that in the development of any new field, there are three basic stages of development. The first stage is the stage of relative chaos. Hoffmeister, who serves as Skandia's Global Director of Intellectual Capital Reporting, has seen this firsthand in his own immediate field, intellectual capital (reporting), for which there presently are no standards. This is the stage, then, in which many competing and discordant ideas are bandied about, even in the absence of commonly held views on what the questions are, or the guiding principles that help determine the answers.

The second stage is the stage in which commonly held principles start to form. In knowledge management, for example, we could say that if the frameworks put forth in the new KM, such as the KLC and the unified theory of knowledge that it embraces, rise to the level of commonly held knowledge, the second stage of KM will have at least begun. *And once it has, and after it has established itself on a more stable footing, it will be time to proceed to the third stage and to consider the development and adoption of standards for KM. But not before!*

We are nowhere near the third stage of development in KM, nor are we barely beyond the first. That scene has yet to be played out. At best, we are in the early period of the second stage, but even that is mere speculation. To undertake the development of standards for KM at this juncture in time borders on the irresponsible. Let's consider some of the errors reflected in many of today's KM standards efforts to help illustrate why current efforts at forming KM standards are highly premature.

First, how can we possibly have standards for KM when most such efforts fail even to make the most rudimentary distinction between knowledge management and knowledge processing? Worse yet is the constant confusion that continues to exist about the two, the result being that some of what we see in related standards efforts are arguably KM-related while others are clearly knowledge processing related. This, of course, is necessarily accompanied by

expressions of *what KM* is that are foggy at best. Is it to enhance knowledge processing? Is it to enhance knowledge sharing? Is it to help harvest the economic value of "intellectual assets"?

A quick survey of the many KM standards efforts now under way will reveal that apparently it's all of the above. Which, of course, means that it's none of them, for how can we have standards for the same thing in different places around the world that in very material ways contradict one another? The answer is we can't, or at least we shouldn't. If ever there was a case of standards-making efforts being launched prematurely, this is it. We're now actually codifying our prematurity in the form of conflicting sets of standards.

Second, it is embarrassing to watch the deeply ironic manner in which standards in KM are being pursued in venues that, themselves, are engaging in KM without the slightest bit of self-consciousness about what it is they are doing. The American National Standards Institute (ANSI), the British Standards Institute (BSI), and the International Standards Organization (ISO) are three that come immediately to mind, although KM standards development efforts are now also under way under the auspices of several other standards-making organizations around the world. But what is a standardsmaking organization if not a formal knowledge management system, itself? And what, for that matter, is a standard if not a set of codified knowledge claims? And what is a KM standard if not a set of knowledge claims about knowledge management or knowledge processing?

Indeed, the processes that organizations like ANSI, BSI, ISO, and others prescribe and enforce are normative knowledge processing systems, the purposes of which are to produce validated knowledge claims (i.e., standards). The irony here is palpable! Let us put it in explicit terms. If a group of people is interested in developing standards for how organizations and their managers should go about managing knowledge processing, why not just adopt the processes already in use by the standards-making organizations themselves, which they (the would-be KM standards developers) apparently already find acceptable? Why not just take the ANSI process, or the ISO process or what have you, embrace it as the KM standard, declare victory, and move on?

It would be even more ironic if any of the now underway KM standards efforts were to actually produce outcomes that conflicted with those enforced by the standards-making bodies under whose auspices their work occurred. The hypocrisy here would be overwhelming. How can a group accept the validity of one set of rules for knowledge processing and yet ratify another, even as it relies on the legitimacy of the first set (which it then implicitly rejects) as the basis for producing the latter? No, intellectual honesty would seem to demand that any KM-standards related group that subordinates itself to the standards-making procedures of an ANSI, BSI, ISO, or what have you should immediately be seen as having embraced, and thereby endorsed, the knowledge management model of their overseers and the knowledge processing systems

they prescribe. Why, then, continue with the development of yet another, new set of KM (or knowledge processing) standards as if the ones practiced and enforced by the overseers didn't already exist?

Third, in light of this criticism of KM standardization efforts, what can we say about the future of such efforts? (1) We can safely assume, we think, that the maturity of standards that conflict with one another in KM will have precisely the opposite effect of what their makers hope for. The only thing worse than conflicting theories and practices in KM is standards that codify them rigidly into place.

(2) It is our hope that the overseers now reigning supreme in the development of KM standards efforts will themselves be seen as enforcers of an approach to KM *and knowledge processing* that is unduly determinative of outcomes and biased in its effects. Why? Because even a cursory inspection of the standards -making processes enforced by ANSI, BSI, ISO, and the rest will reveal the fact that they are (a) KM organizations themselves whose methods we are expected to adopt, and (b) whose enforced knowledge processing system is fundamentally communitarian and therefore deeply flawed.

Communitarian systems employ consensus as the basis of knowledge claim evaluation in their knowledge processing schemes. According to that logic, a standard could (and will) be adopted simply by virtue of its popularity or the degree to which a consensus builds behind it. Once this consensus is reached, communitarian systems are not open to continuous testing and evaluation of the knowledge claims backed by the consensus, for the simple reason that the criterion of truth in such systems is the consensus that has already been reached. Never mind the veracity of the claims themselves or their closeness to the truth. Communitarian systems seek harmony and agreement amongst their stakeholders, regardless of the quality or integrity of the knowledge they produce.

Let us conclude our remarks on this subject, and the book itself, then, by pledging to do whatever we can to continue to call attention to the prematurity and hypocrisy of KM standards-making efforts under current industry and institutional circumstances. This, too, is liable to change under the influence of the new KM in the days ahead. •