
The Eleven Deadliest Sins of Knowledge Management

Liam Fahey
Laurence Prusak

A core tenet of any organizational learning project is that without detecting and correcting errors in “what we know” and “how we learn,” an organization’s knowledge deteriorates, becomes obsolete, and can result in “bad” decisions. Because systematic attention to knowledge management is relatively recent, it is particularly important to detect these errors so that knowledge management does not become yet another management fad that promised much but delivered little. If we do not identify and try to resolve these errors, “what we know” about knowledge management may become little else but mythology. As a consequence, we will be faced with the ultimate knowledge irony: efforts to manage knowledge are themselves based upon faulty knowledge principles.

The purpose of this article is to draw attention to a set of pervasive knowledge management errors. These reflections are based on the authors’ observing or partaking in over one hundred knowledge projects over the past five years or so. The focus is on fundamental errors, that is, errors that if left uncorrected inhibit genuine knowledge from being developed and leveraged. These are errors associated with the concept of knowledge itself: how knowledge is understood in organizational settings and how that understanding impedes knowledge management.

Error 1: Not Developing a Working Definition of Knowledge

If knowledge is not something that is different from data or information, then there is nothing new or interesting in knowledge management. Yet many managers seem determinedly reluctant to distinguish between data and

information on the one hand and knowledge on the other; and, more importantly, they seem reluctant to consider the implications of these distinctions.

The tendency to avoid grappling with what knowledge is should not be surprising. There is little in the education, training, or organizational experience of managers that prepares them for the deep-seated reflection and understanding required by the concept of knowledge. Moreover, this situation is exacerbated by some recent popular management literature that directly advocates not making distinctions between these concepts. The argument advanced by these authors is that contemplation of such distinctions distracts managers from the necessary task of managing.¹ However, reflection upon concepts and the distinctions among and between them is the essence of the process of "knowing" or learning.

This is a critical error. It contributes directly to all of the errors noted below. Also, avoidance of grappling with a working understanding of knowledge leads to a dysfunctional environment for knowledge work. Many executives have told us they were extremely reluctant to even use the knowledge word and that they felt the anti-knowledge culture of their organizations compelled them to do knowledge work by stealth. "We had to disguise our knowledge project within a data warehousing architecture plan" is a true and representative response. In fairness, firms have been assaulted, at least since the 1960s, with multitudes of theories and nostrums that have often proved to be of questionable value. This has made many executives skeptical, if not downright hostile, to new ideas and programs.

Error 2: Emphasizing Knowledge Stock to the Detriment of Knowledge Flow

When knowledge is equated with information, it should not be a surprise to find it defined principally as a stock rather than as a flow. It is viewed as a thing or object that exists on its own, that can be captured, transmitted among individuals, and stored in multiple ways within the organization. Indisputably, this "stock" perspective tends to dominate organizations' thinking about knowledge. This has come about in part because several early examples of knowledge "success" focus on articulated and documented stocks of knowledge such as Dow's work on patent values and McKinsey's rapid-response system.

The notion of flow, however, suggests a radically different conception of knowledge. It is in constant flux and change. It is central to day-to-day doing and being. Individuals create it and it is largely self-generating. Moreover, it connects, binds, and involves individuals. In short, it is inseparable from the individuals who develop, transmit, and leverage it.

The prevalent view of knowledge as stock is grounded in large measure in the thrust of every educational system from grade school through university: learn the facts and regurgitate them as required in the relevant examination.

This orientation is in turn reflected in and reinforced by the pervasive information technology approach to the management of data and information: capture, store, retrieve, and transmit. Although organizations obviously need to manage their data and information using these technology-centered models, knowledge is a substantially different thing and thus needs different models.

The implications for how organizations approach all facets of knowledge management are profound. In many firms, knowledge simply becomes another object to be managed. It is viewed as something separate from the organizational processes that help generate and nurture it. Not surprisingly, therefore, managers all too often do not see themselves as part of the knowledge process, but rather see it as happening outside of them.

Error 3: Viewing Knowledge as Existing Predominantly Outside the Heads of Individuals

Implicit in the observations above is that any discussion of knowledge is meaningless in the absence of a "knower." Knowledge is what a knower knows; there is no knowledge without someone knowing it. Knowledge therefore must be viewed as originating "between the ears" of individuals. Taken literally, the need for a knower raises profound questions as to whether and how knowledge can exist outside the heads of individuals. Although knowledge can be represented in and often embedded in organizational processes, routines, and networks, and sometimes in document repositories, it cannot truly originate outside the heads of individuals.² Nor is it ever complete outside of an individual.

In this view, knowledge is shaped by (among other things) one's initial stock of knowledge, what goes on inside one's head (that is, how one reasons), and the inflow of new stimuli (such as new data and information). Flow, therefore, is also central to what happens between one's ears.

Yet organizations seem to view knowledge as if it has a life of its own. They dub strikingly mundane databases as "knowledge bases," they talk of search engines as if they were human brains, and they extol executive expert systems as if the human mind were incidental to their construction and use. This attempt to dress up decades old technologies and concepts in new "knowledge" clothing is one of the more serious distractions faced by knowledge advocates.

In conjunction with an emphasis upon knowledge as stock, this error reinforces organizational tendencies to manage and massage ever more complex and interconnected databases and to construct even more elaborate information structures. This would not be so bad except that it shifts the focus of knowledge and knowledge work away from individuals—without whom knowledge can be neither generated, transmitted, nor used.

Error 4: Not Understanding that a Fundamental Intermediate Purpose of Managing Knowledge Is to Create Shared Context

If knowledge exists ultimately within individuals, and it is individuals participating simultaneously in multiple group processes who make and execute key decisions, then a fundamental purpose of "managing knowledge" must be to build some degree of shared context. "Shared context" means a shared understanding of an organization's external and internal worlds and how these worlds are connected. Shared context is dynamic: knowledge as flow implies that any shared understanding is likely to change over time, and sometimes may do so suddenly. In the absence of shared context, individuals' differing perspectives, beliefs, assumptions, and views of the future are most likely to collide and thus immobilize decision making.

Yet many organizations still appear so caught up in the model of knowledge as stock that explicating, challenging, and aligning distinct views receives little systematic attention. In short, individuals' understanding of the world around them—perhaps the most critical link between knowledge and decision making—is ignored. This error obviously stems from a "stock" view of knowledge, but its roots lie deep in the prevailing cultures of most organizations. Establishing, challenging, and aligning shared context requires decision makers to engage in open, honest, supportive (and yet critical), and reflective dialogue.³ However, knowledge is a direct outcome of experiences, reflection, and dialogue—three activities that use up that most precious managerial asset: namely, time. Few firms feel they can afford to budget directly for these activities, yet little knowledge is ever developed without them.

A disregard for shared context means that the generation, transmission, and use of knowledge is not seen as an activity that brings individuals to deeper understanding through dialogue. As a result, information remains simply a pattern of disjointed and ill-structured data points or events. Without such dialogue the path from information to knowledge is difficult to traverse.

Error 5: Paying Little Heed to the Role and Importance of Tacit Knowledge

A "head centered" (or perhaps, more accurately, an *embodied*) view recognizes the central role of tacit knowledge in shaping and influencing explicit knowledge. Tacit knowledge entails a body of perspectives (e.g., our view of customers is framed by our firm's experience in North America), perceptions (customers seem disinclined to try our new product), beliefs (investment in new technology will lead to breakthrough new products that will create new customer needs), and values (do what is right for the customer). Tacit knowledge is the means by which explicit knowledge is captured, assimilated, created, and disseminated.

Yet in spite of the emphasis upon tacit knowledge in both traditional epistemology and the recent knowledge management literature, organizations seem especially reluctant to grapple with its management. One unfortunate consequence is that in redoubling their commitment to managing explicit knowledge, organizations discover to their dismay that their efforts are thwarted by the very phenomenon they choose to downplay or ignore. As more emphasis is placed upon refining and extending customer satisfaction surveys, managers' perceptions, projections, and values intervene to preclude any genuinely new insights into customers' behavior.

The most profound reason for this error is that managers simply do not understand the nature of tacit knowledge, its attributes, or its consequences. Thus, although they admit it is "there," they fear it is inaccessible and impossible to influence. Thus, their own tacit knowledge about knowledge serves to limit their understanding of the real nature of knowledge—both tacit and explicit.

If this error persists, the development and leveraging of explicit knowledge is largely stifled. Consider the following case. One organization considered its extensive services to be the dominant reason its customers continued to do more business with it. In in-depth interviews with many of these customers, however, service was ranked as only the fifth or sixth most important purchase criterion. Yet, key members of the management team not only refused to believe the interview findings, they steadfastly refused to entertain the possibility that their long-held semi-tacit assumption might be wrong.

Error 6: Disentangling Knowledge from Its Uses

Knowledge is about imbuing data and information with decision- and action-relevant meaning. This is the vital role of human intervention.⁴ Information about customers becomes knowledge when decision makers determine how to take advantage of the information. In this way, knowledge is inseparable from thinking and acting (see Error 7).

Yet many organizations disconnect knowledge from its uses. A major manifestation of this error is that so-called knowledge initiatives, projects, and programs become ends in themselves. Data warehousing, customer satisfaction surveys, and industry scenarios degenerate respectively into technological challenges, management games, and clashes among proponents of different scenario methodologies. Their relevance for decisions and actions gets lost in the turmoil spawned by debates about appropriate data structures, best survey designs, and alternative techniques for imagining specific industry futures.

This error arises directly from a number of decisively false assumptions in the way many organizations approach knowledge management. First, access to information is not equivalent to insight, value, or utility. Examples of managers recognizing in retrospect how they should have derived insight from particular data and information are legend in every company. Second, the value of data

and information is often anything but obvious. Sometimes it is only after considerable discussion and dialogue that the decision relevance and usefulness of data and information becomes evident. Discerning the appropriate marketing strategy responses to new requests of key customers often requires extensive analysis and detailed projections of the forces shaping customers' behaviors. Third, a common tendency to firmly segregate knowledge users ("decision makers") from many of those involved in generating knowledge further serves to separate knowledge from its potential uses. The universal use of the term "knowledge worker," as distinct from workers who presumably don't have or use knowledge, is a prime indicator of how common this sort of error is. A recent survey by Mark Fruin on how knowledge is understood and valued at Toshiba points out the fatuousness of these labels.⁵

A critical implication of this error is that the knowledge efforts of many organizations are misdirected. In short, they frequently commit extensive resources and time to refining and perfecting data and information at the expense of deriving decision and action implications. This tendency is vividly manifest in the extraordinary lengths to which many organizations go to ensure that their customer surveys meet every statistical standard, that unbiased questions are posed, and that the data collection process does not influence the results. Only later is it discovered that the data generated is not terribly helpful in many critical marketing decisions and actions (such as the design of new products, changes to current marketing strategy, or development of rapid responses to the actions of competitors).

Error 7: Downplaying Thinking and Reasoning

Knowledge generation and use at the level of individuals and groups is a never-ending work-in-progress. At its core, however, getting to different states of knowledge development requires some form of reasoning. For example, a sequence of observations about how customers use a product may lead to insights about desired product modifications, potential new customer solutions, or ways in which existing products might be better customized to specific customers' needs. Explicating thinking and reasoning processes is especially critical in the case of explicit knowledge.

Yet it is always a shock to observe how little attention is paid by allegedly well-managed organizations to their modes of reasoning: the nature of the analytics inherent in points of view or arguments; the assumptions underlying particular models and metaphors; or the relationship between a mode of reasoning and its "logical" outcomes and consequences. Organizations are thus often unaware why changes occur in a particular stock of explicit knowledge. Hence, they are unable to test the rationales for or validity of such changes.

While there are many organizational causes of this error (e.g., an organization's culture does not tolerate articulation and consideration of conflicts in reasoning), the dominance and pervasiveness of tacit knowledge is one of its

principal direct sources. Managers' deeply held, widely shared, and largely untested perceptions of and assumptions about customers' changing behaviors overwhelm data describing customers' responses to a competitor's new product introduction. The following words might be heard: "Customers will quickly see how inferior the competitor's product is to ours. As soon as they do, they will return to our fold." What goes undeveloped and untested is a set of alternative explanations of the documented customer behaviors and their implications for the organization's marketing strategy.

The obvious implication is that both tacit and explicit knowledge solidify and ossify. Unless distinct modes of reasoning—such as alternative explanations of customers' responses to a competitor's new product introduction—are articulated and assessed, radical disjunctures in knowledge content or breakthroughs in insight are considerably less likely to emerge. In short, unless the "frames"—points of view embodied in perceptions, beliefs, assumptions, and projections about the future—are broken by challenging prevailing modes of thinking and reasoning, knowledge generation and use will be severely restricted.

Error 8: Focusing on the Past and the Present and Not the Future

If the intent of knowledge is to inform and influence decision making, then its focus must be on the future. Although we cannot know the future, every strategy, decision, and action is, by definition, premised upon some view of the future. Knowledge, as distinguished from raw data and information, can create a shared context for organizational members to address the future.

Yet in most organizations, knowledge is still predominantly used for understanding past and present change. The implications of such change for future decision making and action receive only secondary attention at best.⁶ Although some organizations are becoming more comfortable in developing "memories of the future" through the use of scenarios and other techniques,⁷ it is still the rare organization that explicitly makes discussing the future the driving focus of its knowledge work. Of course, it hardly needs to be noted that database and information management generally devote little explicit attention to consideration of the future.

A number of causes for this error warrant particular attention. First is the greater comfort and ease with which individuals are able to collect data and generate information about the past and the present. Creating explicit projections unavoidably entails making judgments about the future. The strong natural tendency to avoid the personal, political, and organizational risks associated with exposing one's thinking and reasoning about the future often conspire to stop development of "future knowledge" right in its tracks. Second is the common misunderstanding of knowledge and the future: the goal is not to know the future, rather it is to know what projections of the future inform management's

thinking about its strategies and actions. Third is a failure to recognize that alternative projections of the future enlighten our understanding of the present.⁸ Moreover, what we take to be knowledge about the present may be largely illusion. As a consequence, the value of our knowledge efforts for decision makers may be far less than we presume.

Error 9: Failing to Recognize the Importance of Experimentation

Experiments are a crucial source of the data and information necessary for the invigoration of knowledge, and in most respects, the creation of new knowledge. Experiments include trying new approaches to analysis, initiating pilot projects, doing things on a trial-and-error basis, and allowing individuals to assume additional tasks and responsibilities. Customer experiments might include directly involving customers in specific stages of product development, asking and supporting customers to use the product in different ways, and testing new ways to deliver solutions to specific customer segments.

Although experiments are a naturally occurring phenomenon in every organizational setting, few organizations explicitly seek to continually create and leverage experiments for knowledge purposes. Indeed, a failure to recognize the potential experimental value of ongoing activities is all too evident in most organizations. For example, many firms now document their "best practices" as a means of encapsulating knowledge at work and disseminate such descriptions and results to a wide range of potential internal users. However, such documentation also tends to suppress any inclination to track and monitor improvements and innovations in each best practice as it is applied and enhanced by various user groups.

The use of technology tends to result in standardized approaches to collecting and structuring data and to transferring information. This tendency is reinforced by command-and-control, hierarchy-driven organizational cultures that specify precisely what individuals can (and cannot) do. The result is an emphasis on exploitation over exploration.⁹ Organizations invest time and resources in improving current modes of data gathering, enhancing the efficiency of IT, calibrating information structures, and involving more individuals in information and knowledge routines. One consequence of this is an emphasis on simply refining and sharpening what we already know. What is downplayed, and often dismissed entirely, is the willingness to explore: to do new things, to do old things in new ways, and to learn from both these activities. Distinctly new knowledge stems from experimenting.

Error 10: Substituting Technological Contact for Human Interface

The veritable explosion of information and communication technologies has created the means to capture and transmit data and information at rates and speeds that were unimaginable merely a few years ago. Data warehousing, search engines, groupware, and newly emerging client-server systems are only the tip of the iceberg. Information technology budgets continue to escalate, vendors are sprouting up like mushrooms, IT professionals are increasingly assuming knowledge titles, and many knowledge projects quite rightly depend upon intensive technology use. There is widespread tendency to validate significant investment in IT by reference to its contribution to developing and leveraging knowledge in new and effective ways.

Unfortunately, one pivotal error underlying some uses of IT severely limits its potential contribution to organizational knowledge: technological contact is equated with face-to-face dialogue. Although IT is a wonderful facilitator of data and information transmission and distribution, it can never substitute for the rich interactivity, communication, and learning that is inherent in dialogue. Knowledge is primarily a function and consequence of the meeting and interaction of minds. Human intervention remains the only source of knowledge generation. This in turn has produced immense frustration on the part of executives. One result of this frustration is further investment in IT, which without the requisite change in the understanding of knowledge, in turn only leads to further frustration and disappointment.

Error 11: Seeking to Develop Direct Measures of Knowledge

A reasonable and sensible question now being raised by many concerned senior managers is: How will we know if our efforts to manage knowledge produce satisfactory results? Or, stated differently, where is the pay-off to knowledge projects?

Regrettably, it seems that an increasing number of organizations seek to measure knowledge directly rather than by its outcomes, activities, and consequences. Thus, they emphasize the scope, depth, number, and quality of databases; the numbers of individuals, units, and departments connected technologically; the number of requests or "hits" pertaining to intranets, key knowledge sources, and information pools; and the number, variety, and extensiveness of knowledge projects or initiatives. Yet a moment's reflection will convince even the most diehard metric devotee that such indicators do not provide any sense of an organization's stock or flow of knowledge or its contribution to decision making and organizational performance.

However, some firms have developed "proxies" for showing the outcomes or consequences of knowledge-based activities. These metrics include patents, new products developed and introduced, customer retention, and process

innovation. Yet it seems fair to suggest that in their thirst for metrics and measures, an embarrassing number of organizations still put the measurement cart before the knowledge horse. This vain pursuit of metrics reinforces many of the errors previously noted. It misconstrues what knowledge is, it consigns human intervention to a secondary role, it further disconnects knowledge from its users and its uses, and stock is given prominence while flow, because it is so difficult to measure, receives minimal attention.

What Can Be Done?

Identifying these errors in one thing; avoiding them is quite another. There are three critical sets of actions that managers can take to avoid these errors and move their organization toward becoming a more knowledge-driven enterprise.

First, managers need to continually reflect on knowledge as an organizational phenomenon. Because knowledge is distinct from data and information, considerable discussion and reflection is required to develop a consensus about knowledge itself. In their quest to demonstrate the results of their knowledge efforts, far too many organizations have neglected to develop a shared understanding of what the knowledge-driven enterprise is all about. Manager need to take the following steps:

- Develop shared understanding at local levels; since knowledge tends to be a "local" phenomenon, it is impossible to do so quickly at the level of the enterprise.
- Allow individuals frequent opportunities to discuss and debate what knowledge is.
- Help individuals identify their current and desired knowledge roles.
- Ask individuals to identify knowledge implications for group behaviors and processes.

Second, managers must be obsessive about noting and correcting errors in their stock of knowledge—or, more precisely, in what they think they know. Such attention must go beyond a mere verification of so-called "facts," that is, descriptions of what is unambiguously true. Unfortunately, this sort of attention is the inevitable focus of a data- or information-dominated view of knowledge. However, facts constitute only a small part of the "stuff" of decision making.

Therefore, managers need to continually expose knowledge content and subject it to scrutiny in every possible way. In doing so, they will quickly come to recognize one critical consequence of the flow perspective of knowledge: substantial portions of knowledge content are always tentative, temporary, and subject to change. There are few absolutes. It is imperative that managers ask the following questions with regard to customers:¹⁰

- What do we know, or think we know, about different aspects of our customers?
- Do we know our own perceptions, beliefs, assumptions, and projections about different categories of customers?
- How does what we know and what we don't about customers affect specific decisions?
- What do we need to know about potential customers to make specific decisions and how is it different from what we think we know?
- What errors might reside in what we think we know?
- What might be the consequences of these errors?
- How might we rectify these errors?

Third, managers must be vigilant about detecting and correcting errors in their processes of knowing—the generating, moving, and leveraging of knowledge throughout the firm. Such learning processes are often deeply embedded in the way the organization does things. Questions managers need to ask are:

- Which individuals play what roles in developing and testing information?
- Which *individuals*, or categories of *individuals*, are not involved in dialogue around specific issues and topics? How might their involvement affect the content and flow of knowledge?
- How is knowledge flow facilitated or impeded by the organization's structure and systems?
- How does tacit knowledge influence the generation and transfer of explicit knowledge?
- How is technology used to unearth and influence tacit knowledge?
- What role do experiments play in knowledge generation?

In sum, an organization must engage in critical, sustained, and honest self-reflection about the errors noted in this article. By doing this, it can avoid the pitfalls that are evident in the approaches of many organizations' attempts to work with knowledge.

Notes

1. For example, see, Thomas A. Stewart, *Intellectual Capital* (New York, NY: Doubleday/Currency, 1997).
2. Our perspective on this controversial subject is strongly influenced by R.R. Nelson and Sidney G. Winter, *An Evolutionary Theory of Economic Change* (Cambridge, MA: Harvard University Press, 1982).
3. By dialogue here we mean not just discussion but the presentation and consideration of different views and perspectives with the purpose of developing a distinctly different or new view or perspective. For further insight into the distinction between discussion and dialogue, see Peter M. Senge, *The Fifth Discipline: The Art and Practice of the Learning Organization* (New York, NY: Doubleday/Currency, 1990).

4. The centrality of the human role in identifying and executing decision and action implications again reinforces the criticality of errors 3, 4, and 5.
5. W. Mark Fruin, *Knowledge Works: Managing Intellectual Capital at Toshiba* (New York, NY: Oxford University Press, 1997).
6. Note here again the particular relevance of error 6.
7. See, for example, the many illustrations of scenarios as a means to develop alternative memories of the future in Liam Fahey and Robert Randall, eds., *Learning from the Future: Competitive Foresight Scenarios* (New York, NY: John Wiley and Sons, 1998).
8. This point has been powerfully argued by Brian March, "Using Scenarios to Identify, Analyze and Manage Uncertainty," in Liam Fahey and Robert Randall, eds., *Learning from the Future: Competitive Foresight Scenarios* (New York, NY: John Wiley and Sons, 1998).
9. This distinction has been made by James G. March, "Exploration and Exploitation in Organizational Learning," *Organization Science*, 2/1 (February 1991).
10. Obviously the same types of questions could be asked of many entities or topics.

Copyright of *California Management Review* is the property of *California Management Review* and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.