

Knowledge Management: An Introduction

Marc Demarest
demarest@hevanet.com

The recent rash of assertions that knowledge is -- somehow or other -- the key to effective competition, marketplace distinction and profitability in the global post-capitalist economy serves to underscore the extent to which the average firm is unsuited for this new economy, in three critical ways.

If it is the case that effective management of any firm's marketplace performance is based primarily on:

- a shared, more-or-less **stylized notion of how value is created, maintained and distributed in within the firm** and shared with upstream, sidestream and downstream elements of the value chain within which the firm participates
- a set of **processes and systems -- technical or human -- that support and help channel the firm's value creating activities** along the lines suggested by the shared, stylized model
- a set of **key metrics that link the value-creation process within the firm to more-or-less objective measures of the firm's success** (stock price growth, P/E ratios, customer satisfaction indices and market share measurements)
- a set of **command-and-control systems¹ -- knowledge management infrastructure -- that monitor the efficiency and effectiveness of that value creation process**, indicate opportunities for performance improvements and generally signal the relative rise or decline in value-creation at various points within the firm

then the "knowledge problem" for the typical firm is this:

¹ It has become commonplace to denigrate the whole notion of command-and-control systems as being anti-empowerment and therefore politically incorrect, not to say destructive of the core values that are alleged to make for competitive advantage in the New Economy. The fact of the matter is that, so long as the Securities and Exchange Commission, the Internal Revenue Service and shareholders exist, command-and-control systems have a valuable role to play in the operation of the firm.

- there is no commonly-held model for knowledge creation and dissemination within the firm
- there are no processes or systems focused on supporting these activities
- there are no metrics for evaluating the role or effectiveness of knowledge-creation and knowledge-dissemination activities
- there are no command-and-control systems focused on measuring and evaluating the various knowledge-creation and dissemination processes in the firm.

The challenge, for a firm whose strategy requires effective participation in knowledge-intensive sectors of the global economy, is embodied in the answer to six key questions:

- what does our culture, and our actions as managers, say about the value of knowledge in the firm, and what do we ourselves believe about the value, purpose and role of knowledge?
- how is knowledge created, embodied, disseminated and used in this firm, and what is the relationship between knowledge and the innovations and performance this firm requires to succeed in its strategic objectives?
- what strategic and material commercial benefits do we expect to gain from more effective knowledge management and the performances created by effective knowledge management?
- where is our firm in terms of the maturity of its knowledge systems?
- how must we organize for knowledge management?
- what role does information technology play in our knowledge management program?

Commercial Knowledge: What Is It, Exactly?

Jorge Luis Borges' essay "The Analytical Language of John Wilkins" contains a passage that has often been quoted in discussions about the nature of knowledge. The salient passage reads in part:

On those remote pages it is written that animals are divided into (a) those that belong to the Emperor, (b) embalmed ones, (c) those that are trained, (d) suckling pigs, (e) mermaids, (f) fabulous ones, (g) stray dogs, (h) those that are included in this classification, (i) those that tremble as if they were mad, (j) innumerable ones, (k) those drawn with a very fine camel's hair brush, (l) others, (m) those that have just broken a flower vase, (n) those that resemble flies from a distance.

Borges is ostensibly quoting a Chinese encyclopedia that Wilkins found fascinating, pointing out that knowledge comes to us in many forms, some of which we are not able to assimilate because of our own assumptions about what knowledge is, and how it is systematized. The example below strikes us as fantastic because (a) it lacks internal symmetry and balance and (b) the grid of distinctions it draws do not correspond to the categories we habitually manipulate in our daily lives.

Yet, it reflects a kind of knowledge that was presumably productive for Chinese court philosophers: presumably, it *worked for whatever purpose it was developed and used*. This is the essential nature of knowledge, as we have to deal with it commercially: good commercial knowledge, valuable knowledge, is knowledge that **works**.² Its truth value is incidental to its ability to generate desirable commercial performances.

Commercial knowledge, I would argue, is different in kind, not degree, from philosophical and scientific knowledge. Scientific knowledge as the paradigmatic form of all knowledge is conventional -- as academics are fond of saying, all disciplines aspire to be physics -- and very wrong in this case. Similarly, philosophical knowledge, under assault from post-structuralist disciplines that argue that truth is embedded in language and therefore inaccessible, can't show us commercial truth, as if understanding the essence of commerce would help us win in a marketplace that is more than half a collectivity of perceptions.

The goal of commercial knowledge is not truth, but effective **performance**: not "what is right" but "what works" or even "what works better" where better is defined in competitive and financial contexts. Partial, or even empirically

² This is a position not too far from American pragmatism (James, Dewey) and the idea of knowledge-for-action as exemplified in the work of Chris Argyris.

incorrect or philosophically muddled, bodies of knowledge have been, and are effective. The Ptolemaic universe worked. Classical mechanics works. Aristotle's notion of spontaneous generation -- the theory that suggests that piles of grain, burlap, dark corners and damp earth breeds rats -- **works** if what you are trying to do is (a) increase rat populations or (b) exterminate rats.

Commerce is about the provisional: about rules-of-thumb, swags, and truths that are highly productive, and then become unproductive overnight.

All commercial knowledge is provisional, partial, muddled -- and, when it is good, it works.³

Some other attributes of commercial knowledge that are counter-intuitive but broadly accurate:

- All commercial knowledge is social: produced and shared among a network of human and non-human actors within the firm (and increasingly across the firm's boundaries).
- All commercial knowledge is traded, by knowledge workers, in knowledge economies that are today operating within every knowledge intensive firm in the world. What is traded for is either other different kinds of knowledge, or status within the knowledge worker guild within the firm.

³ Commercial knowledge is very close to what the French call *bricolage*: the provisional construction of a messy set of rules, tools and guidelines that produce according to the expertise and sensitivity of the craftsman, not the empirical accuracy of the rules, tools and guidelines.

Modeling Knowledge Economies Within The Firm

I have implicitly suggested, in an earlier section of this essay, that:

- all firms have knowledge economies operating within them already
- all firms' knowledge economies operate in sufficiently similar ways to justify generalizations about how knowledge is produced and consumed in commercial firms.

Schematically, the knowledge economy within a firm looks something like this:

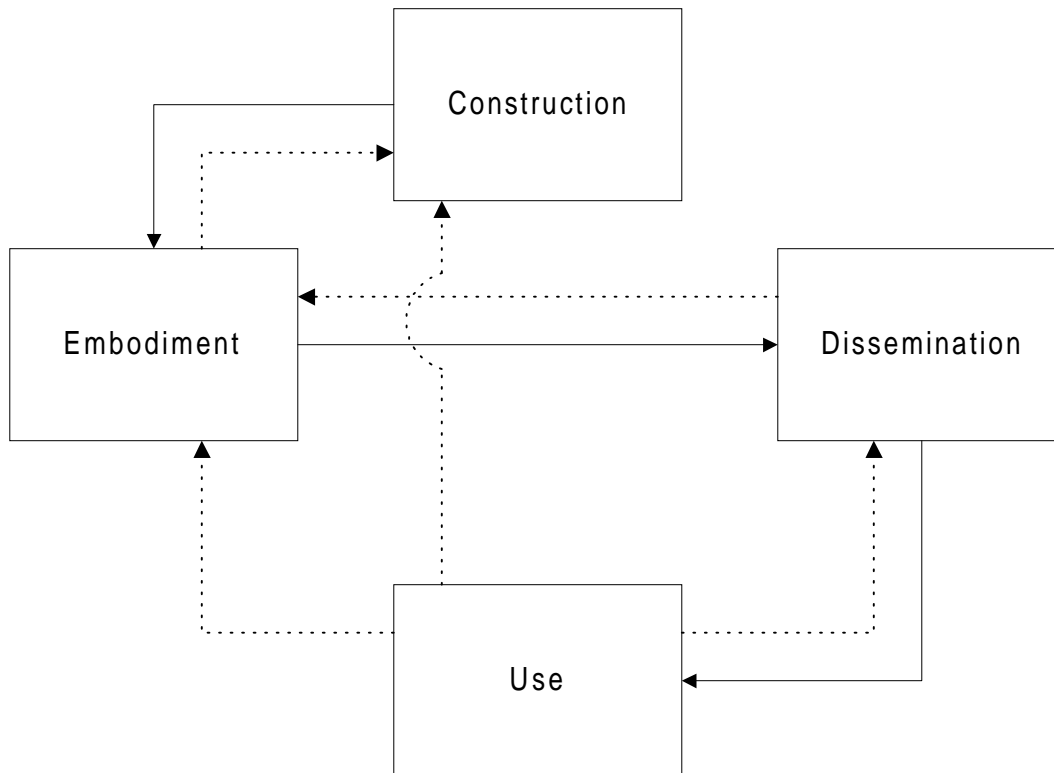


Figure 1 -- The Knowledge Management Process⁴

⁴ For more information on the background to this model, refer to Clark and Staunton, pp.51-77 passim.

In the context of this model, and the generic commercial knowledge economy:

- *construction* refers to, in this case, the process of discovering or structuring a kind of knowledge: how to sell a particular product to a particular market, for example, or how to diagnose a particular kind of customer problem.
- *embodiment* refers to the process of choosing a container for knowledge once it is constructed. As we suggested earlier, the container for much of the firm's knowledge is selected by the intelligent desktop, and **that container is most often a document.**
- *dissemination* refers to the human processes and technical infrastructure that make embodied knowledge -- in this case, documents -- available to the people within the firm that use the documents and the bodies of knowledge those documents contain to perform: to do work for customers, suppliers, business partners and the firm itself.
- *use* refers to the ultimate objective of any knowledge management system: the production of commercial value for the customer.

Intuitively, this is how knowledge gets produced within any firm, and all participants recognize the model as essentially accurate. The fact that this model is detected by researchers into organizational innovation only underscores its essential accuracy.

It is worth noting that the model assumes a certain amount of simultaneity, along a number of vectors:

1. *construction-to-use*: the constructor, and frequently other members of the constructor's circle, typically put knowledge into practice while it is being constructed; indeed, testing-for-use-value is properly a part of the knowledge construction process
2. *construction-to-dissemination*: in a properly functioning knowledge management environment, constructed knowledge is disseminated in a less-than-embodied for testing, validation and critique (testing-for-use-value, or what we know as "review cycles")
3. *construction-through-embodiment-through-dissemination-through-use*: the more formal process is ideally executed only after (1) and (2) above (that is, after both the constructor and a small network of others have performed testing-for-use-value)

Understanding how these processes function in an organizational unit (of whatever size), making judged decisions about the optimal methods for each process (as well as vehicles for embodiment and dissemination), and finally

determining a minimal, optimally effective set of measurements for the entire constellation of processes is a key -- if not the key -- aspect of successful knowledge management.

Knowledge Management: Some Definitions

A Definition Of Commercial Knowledge

Commercial knowledge, for our purposes, should be thought of as:

an explicitly developed and managed network of imperatives, patterns, rules and scripts, embodied in some aspect of the firm, and distributed throughout the firm, that creates marketplace performances.

Explicitly Developed And Managed

It would be possible to argue that all firms produce and use significant knowledge, whether they know it or not, and that therefore explicit development and management of knowledge are irrelevant: knowledge will get created, one way or another.

This is true -- but the issue is not ultimately whether knowledge gets constructed or not. In some sense, the construction of knowledge is encoded in our genes and mandated by our environment; we must at some level get smart or die.

The issue is that firms do not value what they do not:

- explicitly support with business processes
- explicitly manage with people and systems
- explicitly measure.

The test of a knowledge management system is ultimately whether it is a by-product of the firm's operations, or an explicit objective of those operations. I would argue that, very soon now, any firm that treats knowledge as a fortuitous by-product of its operations (treating knowledge in effect as a happy effluvium of doing), rather than an objective of those operations, will be out of business, not because it could not construct knowledge, but because it could not deliver the knowledge created to the marketplace in terms of market-valued performances.

A Network Of Imperatives, Patterns, Rules And Scripts

Knowledge is actually knowledges: plural, multiple networks of knowledge. By network, I mean something very specific: the idea that 'nodes' of knowledge (each node an imperative, pattern, rule or script) are held together by 'links' that specific relationships among those nodes: subordination, agreement,

contradiction, ambivalence, or what have you. Both MindMapping⁵ and Fuzzy Cognitive Models⁶ (FCMs) are attempts to describe visually the ways in which knowledge actually forms itself: as networks.

Knowledge is actually knowledges: many firms are simultaneously supporting networks of knowledge at odds with one another in one or more ways. This is the origin of organizational issues like the failed installation of strategies and directives, the maintenance and pursuit of emergent strategies at odds with the firm's formal strategy, and various kinds of cultural malaise and toxicity.

We might do well to distinguish between knowledge, as a general term describing all of these networks, and *discipline*, a term for any single network.

I would argue that generically the knowledge created in commercial firms falls into four categories:

- *imperatives*: behavioral directives that are unchallenged because they are derived from the firm's dogma -- its strategy, goals and operating plans. This is one area in which commercial knowledge differs substantially from other kinds of knowledge. Where, for example, biology rests on certain fundamental postulates that are at least indirectly demonstrable and therefore "true", commercial knowledge may in some cases rest on empirically untenable, even "wrong" imperatives that are expressions of the firm's overweening need to penetrate a new market, recapture competitive distinction, achieve a certain product gross margin or merely survive. Science cannot use "must"; commercial firms do so all the time.
- *patterns*: predictive models that have a certain longevity, durability and level of universality. Patterns describe the likely shape of scenes that call for particular kinds of knowledge, as in: *if the customer has never called with a product complaint before, you can expect the following kinds of behavior...* Patterns also exhibit themselves as usage models (how customers make use of products or services), buying or procurement behavior, social styles, and so forth. In general, they are generalized, stylized models of behavioral contexts in which a kind of knowledge is to be practiced or deployed.
- *rules*: algorithms and heuristic logic models that define a basic set of guidelines for performing in particular environments. These rules may be algorithmic (that is to say, open to only one interpretation and susceptible to coding in one form or another) or they may be fuzzy; that is, they deal with qualitative evaluations and categories of language, and they are incomplete and require inventiveness on the part of the performer. A model for conducting a BPR workshop constitutes a set of rules, as does a manual

⁵ A popular form on note-taking much used in the United Kingdom.

⁶ A visual description technique described by Bart Kosko in his work on fuzzy logic systems.

describing the administration tasks associated with the diagnosis of a computer system, or a set of guidelines for salesmen on how to talk to CEOs based on the position of their firms in their chosen markets. It was this particular kind of knowledge that early AI proponents sought to capture in so-called “knowledge bases” or expert systems: the rules that, for example, allow senior factory hands to diagnose a piece of machinery on the factory floor by smell, sound or the relative position of a dozen dials, or the skills that allow the firm’s best salespeople to determine with great accuracy during the first sales call whether the prospect will ultimately do business with the firm or not. Rules often do not stand up to empirical scrutiny, because they are often based on collective or individual experience, and therefore seem intuitive or even supra-rational.

- *scripts*: scripts are more than rules; they are prescriptions for performance. Scripts are most often found as recipes, post-mortem examinations, so-called “success stories” or case studies, specific diagnostic tools, skills training (as is the case most obviously with sales training), or methodologies. The salient difference between rules and scripts is that rules sometimes exist in isolation from one another as general postulates-- “if the customer objects to this point, you will not close the sale” -- where scripts are coordinated sets of rules targeted at particular contexts. Scripts give way, in the practice of any knowledge worker, to the employment of heuristics; scripts are, if you like, the raw material of the apprentice knowledge worker, structuring behavior until the knowledge worker can be self-structuring.

Embodied In Some Aspect Of The Firm

Generically, commercial knowledge is either tacit, shared or embodied.

Tacit knowledge is held by a single individual, in the silence of her contemplation or in her particular performances: “top performers” in firms are more often than not people holding highly productive tacit knowledge (an issue we will return to later.) There is some reason to assert that tacit knowledge is not knowledge at all, that knowledge exists only within a practicing network, and that, whatever is held in the heads of “really smart people,” it isn’t knowledge, precisely because it is closely held.

Shared knowledge is held by a group or team, either consciously or in their particular group and individual performances; ‘high-performance teams’ like top individual performers, are typically using shared knowledge to boost their perceived performance relative to other teams in their organizations or elsewhere.

Both tacit and shared knowledge are unembodied: floating, held in memory or in the day-to-day business practices of a small number of people but not formally encoded or available for dissemination or emulation.⁷

When knowledge is embodied -- made formal, explicit and distributable -- it is typically embodied in one of four kinds of things:

- *raw materials, products and services*: knowledge is explicitly used as a component in the development of products or services, and the value that it adds to the finished product or service is typically the basis for the value proposition for that product or service. Consulting firms routinely use information as raw materials. These products may be internal artifacts (that is to say, used only within the firm itself) or actual components of the product or service sold to the firm's customers.
- *machinery and mechanism*: knowledge is embodied in "hard" machinery or mechanisms that change the way the firm's electromechanical infrastructure operates.
- *business practices and processes*: knowledge is embodied in "firm" business processes or practices that change the way the firm's sociotechnical infrastructure operates.
- *environment and culture*: knowledge is embodied in "soft" cultural or organizational values ("open door" policies, for example) that change the way the firm's cultural infrastructure works.

It is possible to think of embodiment in a different way, as a state of knowledge "packaging" along a hard-soft continuum. For example, knowledge of the appropriate way to allocate interrupts to CPUs in a symmetric multiprocessing environment may, once it is no longer tacit or shared and becomes embodied, appear as something as soft as a talk given at a conference, as a research paper published in the proceedings of that conference, as a low-level operating system kernel primitive or as an ASIC. Each embodiment is increasingly hard -- durable and inflexible, long-lived and (increasingly) impervious to reuse. It is tempting to speculate that:

- knowledge moves, over time, from softer to harder forms of embodiment, as a general rule of its behavior
- harder forms are more appropriate to older, commodifying kinds of knowledge.

Distributed Throughout The Firm

⁷ This, I would argue, is why attempts to "clone" the work processes and behaviors of notable high-performance teams and distribute them to other teams within an organization so often fail -- the high-performance team is (tacitly or otherwise) aware that it holds special knowledge, and resists its embodiment and distribution.

For knowledge to translate into some kind of useful commercial performance, knowledge must be distributed from the locale of construction into other parts of an organization or value chains -- ideally to the points in the organization where performance is possible (performances are not equally possible in all locales, although we could wish it to be so). In some cases, embodiment is a prelude to distribution to customers (the classic case being the preparation and printing of technical documentation), while in other cases embodiment (for example in internal training materials or a slide presentation) is the prelude to internal distribution within an actor in the value chain.

Creating Marketplace Performances

As mentioned in the previous section, **the only reasonable purpose of an organized knowledge management practice in the firm is to increase the quality and quantity of marketplace performances**: to enable the firm to sell more and sell better, to support more and support better, to create and keep more, better, customers. A firm's accumulation of knowledge -- particularly under today's models, where for the most part that knowledge remains tacit, unembodied and undistributed -- may ultimately serve some economic good, in terms of making the firm a more desirable acquisition target, or creating a highly-desirable working environment, but the cases of Bellcore, Xerox PARC⁸ and other "think tanks" seems to indicate pretty clearly that this kind of knowledge management does not lead to marketplace performance at all, let alone to dominating marketplace performances that parlay knowledge into revenue or defensible market position.

This implies that the only top-level metrics for a knowledge management system are metrics that ultimately matter are economic ones: market share, revenue, gross margin, concept-to-cash-flow cycle time, customer satisfaction and other commercial success metrics. Other metrics -- concerned with the productivity and efficiency of the knowledge management process itself -- have value to knowledge workers and knowledge managers, but an optimally efficient, high-volume knowledge system that does not lead to dominating market performances has no ultimate value to the firm.

⁸PARC's knowledge leadership in graphical user interfaces, human factors design and object-oriented systems are all compelling examples of this phenomenon. If it is true that nearly every pathbreaking technology has goes through three phases -- innovation, popularization and dominance -- and if it is true that money is made only when the technology penetrates the marketplace in a significant way then PARC can lay claim to being the innovator in all of these technology areas, but can lay no claim to being the popularizer or the firm that benefited from the dominance of any of these technologies. By contrast, 3M has discovered a way of doing business that allows them to both innovate and to cash in, but -- if publicly-available case studies are to be believed -- still relies on serendipity for innovation.

Commercial Knowledge Management

Given this definition of commercial knowledge, a process model for commercial knowledge management would see straightforward. But it isn't, in my experience at least.

What Is Knowledge Management?

In a preliminary way, let me suggest that knowledge management is the systematic:

- underpinning
- observation
- instrumentation
- and optimization

of the firm's knowledge economies.

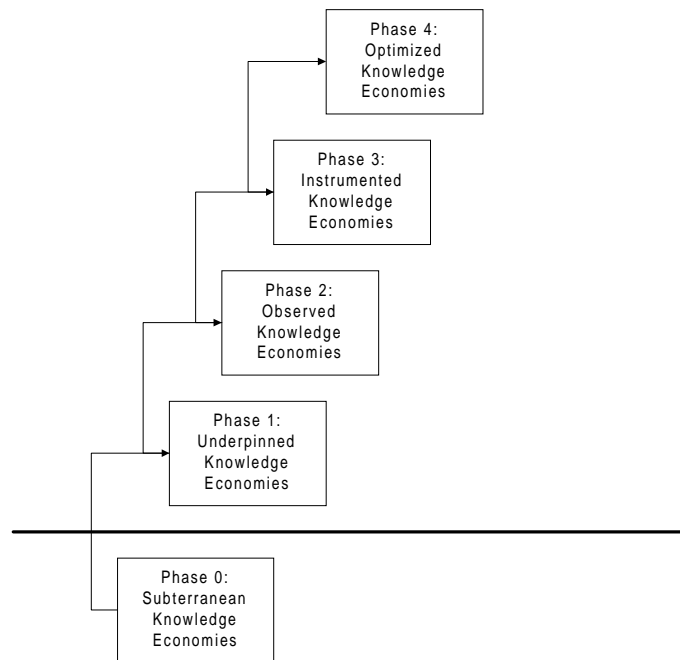


Figure 2 -- Knowledge Management Definitional Schematic

These activities – or perhaps process areas – are sequential. That is to say, it is not possible to optimize a firm's knowledge economies until they are instrumented; instrumentation is selective measurement based on empirical observation and observation of the for-the-most-part underground knowledge economies within the firm requires that these economies be brought out into the light of day, made official, and formally underpinned with technology, process and policy that makes construction, embodiment, dissemination and use of knowledge officially-sanctioned and formally-valued activities within the firm.

This definition suggests what knowledge management might be, or more precisely what people who call themselves knowledge managers might do. What it does not provide is:

- any clear sense of what objectives knowledge managers have in view when they begin the process of underpinning, observing, instrumenting and optimizing their firm's knowledge economies.
- any formal definition of the limits to effective knowledge management within each phase, the dangers inherent in the tasks at hand, or the complexities associated with design and implementation decisions and trade-offs.

The first problem – lack of a clear sense of objectives – is resolved, at least conceptually, with ease: all knowledge management programs ought to be targeted directly at the firm's income statement: at revenue enhancement, cost reduction, or the management of risk associated with marketplace and financial performance. In practice, this is a difficult task, implying a thorough understanding of the firm's internal value chain and the value systems in which it participates, and further implying a formal understanding, at a strategic level within the firm, of the role knowledge plays in the firm's quest to add value.⁹

The second problem cannot, I would argue, be resolved at all except generally without a specific firm's problems against which to apply the general model and definition. Ultimately, all knowledge management programs and problems are local: unique to a particular firm and its knowledge economies.

How To Do It?

So, the problem for commercial knowledge management, in theoretical space, is not "how do we know what we know?" or "what do we do when we perceive?"¹⁰ or "is our knowledge right?" The fundamental questions are:

⁹ The topic of knowledge management strategy – planning and analysis for knowledge management – is woefully underrepresented in the academic and popular literature. We cover this topic elsewhere in the Sequent KMG White Paper series.

¹⁰ Actually, these problems become significant for knowledge managers during and after the transition to Phase 2 (instrumentation).

- does our knowledge work? does it produce economically-valuable performances?
- how well does it work relative to other similar knowledges with which it competes in commercial settings?
- how do we break down the prejudices and assumptive barriers that preclude management of knowledge in the firm?
- how do we create it, practically speaking, as a matter of business practice?
- how do we embody it in products, services, machinery, processes, and cultures?
- how do we maintain and enhance our knowledges?
- when do we scrap them in favor of newer knowledges?

In short, commercial knowledge is valuable in proportion to its productivity in commercial settings: ultimately, in how much leverage knowledge offers us in creating or enhancing customer relationships and their associated revenue streams, in achieving operational efficiencies (through outright cost reductions or process refinements), and in creating the optimal products and services for markets we serve.

This doesn't mean we don't need a theory of commercial knowledge -- we do. Since the fact is that most firms have no idea how knowledge enters, is transformed within, or exits the firm, any theory -- again, no matter how empirically inaccurate -- is valuable, since it:

- gives knowledge managers a model for understanding what would otherwise be impenetrable chaos
- gives knowledge workers a model and a vocabulary for understanding their roles within the knowledge system of the firm.

In other words, knowledge management is to some extent a problem in self-reflexivity -- the thing we are describing does not exist until we describe it, at which point it will probably take on exactly the process attributes we ascribe to it, and then, instantaneously begin to deviate from that description in perceivable, describable ways.

But a theory in and of itself does us little good, as Hammer and Champy have so clearly demonstrated. To create robust bodies of knowledge that work for the

firm, we have to have methods and practices -- ways of working with and for knowledge -- that cover what seem to be the major areas any theory of knowledge management is going to “discover” at work in the firm:

- construction: the “making” of knowledge through complex processes involving creation, theft, traduction and reinterpretation.
- embodiment: the transformation of a knowledge that is tacit -- held in the head of a knowledge worker, shared in relative secret among a small team of confederates -- into processes and practices, machinery, materials and cultures.
- dissemination: the distribution of embodied knowledge throughout a firm or value chain.
- use: the application of disseminated, embodied knowledge to particular problems and classes of problems, with a view to the production of results: to making knowledge work.
- management: the monitoring, measurement and intervention in construction, embodiment, dissemination and use by knowledge managers: people with firms who are responsible for the “knowledge machine” as opposed to its products.¹¹

Once a theory has organized the chaos of knowledge in the firm, and allowed us to develop methods and practices to systematize construction, embodiment, dissemination, use and management, we face the eternal problem of infrastructure: durable cultural and technical systems that facilitate the day-to-day process of knowledge-building. Broadly, there are three classes of infrastructural concerns we have to address:

- cultural infrastructure: today, what one knows is considered central to job security and compensation, which accrue in proportion to the extent to which “what I know” remains in my head. Knowledge management rewards embodiment and dissemination and incessant knowledge creation, and so requires fundamental changes to deep-seated cultural givens in most firms.
- operational infrastructure: human resources organizational practice manages (one would like to say *nurture* instead) human bodies, human feelings, and

¹¹ It's worth noting that “management” in this context is not conventional management of human assets: knowledge managers add tangible value. The primary role of the knowledge manager is to measure and evaluate the productivity of knowledge in the commercial arena: to judge the extent to which a knowledge works in ways the firm wants that knowledge to work, and that the knowledge works better than competitive knowledges do. Thus, the knowledge manager in this model is not a supervisor, standing to one side of the productive process and -- in the service of capital and the owners -- managing the productive mechanism. The knowledge manager defines both the object of knowledge and its effectiveness, and is probably in many cases also the primary constructor or embodying agent of that knowledge.

human perceptions, but not the human mind. Everything about a traditional human resources organization, and much of conventional organizational design and operations, require retooling in a knowledge-centered firm. Job ladders, compensation, work locale, command-and-control structures -- all have to be focused on the individual knowledge worker and her team, and all have to be tied directly to the knowledge-producing capacity of that team.

- technical infrastructure: in the modern firm, knowledge travels through information technology. Later, it will become useful to distinguish between data, information and knowledge -- for the time being, it is broadly correct to say that information technology's primary strategic objective in a knowledge-intensive firm is to facilitate knowledge creation, embodiment, dissemination, use and management, and that each of these five activities requires substantially different complexes of information technology to do its work.

Infrastructure Gets Us To Commercial Performance

Deploying the three kinds of infrastructure properly and completely gets us the knowledge-rich firm.¹² In the same way that TQM disciplines emphasize closing the loop between practice, measurement of practice, and improvement of practice, knowledge management produces financial and competitive results only to the extent that the knowledge management infrastructure is installed, operational and capable of informing on itself -- linking the construction of knowledge to its embodiment to its dissemination to its use to its performance: the productive effects of superior knowledge in the marketplace.

¹² In fact, this allows us to see the "virtual corporation" as nothing more than the collection of knowledge workers who are linked together via an infrastructure under the control of a single collection of equity agents. The firm IS the knowledge management infrastructure it puts into play.

Managing Commercial Knowledge For Economic Value

There is only one reason to practice commercial knowledge management: to enable the firm to sell more and sell better, to support more and support better, to create and keep more, better, customers. If knowledge management doesn't support that objective, it is at best a soft discipline -- useful for enhancing the corporate culture, but finally a nice-to-have, rather than a necessary practice.

Knowledge management pundits argue that in fact knowledge is the capital of the '90s and that, just as the CFO closely manages the capital of the firm, reporting on cash flows and producing balance sheets, so a new corporate officer, the Chief Knowledge Officer (CKO) should be charged with managing the firm's knowledge assets along similar lines. But what is the economic justification for such a position? Does good knowledge management really deliver economic benefit to the firm along the same lines that smart financial management does?

Knowledge Management And Innovation

Clearly, the most obvious link between knowledge management and enhanced economic performance is in the area of innovation. All studies of innovation in the last twenty years have come to pretty much the same conclusion: innovation begins with the construction of a new kind of knowledge within the firm. Sometimes -- as in the classic case of the Post-It Note at 3M -- the knowledge in question is targeted at uses different from those it ultimately serves. Sometimes, as in Hammer & Champy's discussion of IBM Credit, the new knowledge is not product-centric, but process-centric: knowledge about a new way of doing something. Many times -- consider the case of the IBM PC -- the firm's own mechanisms, prejudices and priorities damage the effectiveness of the knowledge at the heart of the innovation.

And, in all cases, we can assume that **innovation will occur without the support of formal knowledge management systems**. So, why bother to formalize knowledge management if innovation will occur anyway? The key is this: as markets become more fragmented and sparse, and market windows open for increasingly brief periods of time, the key to economic viability is not innovation per se, but **repeated innovation with increasingly high levels of reuse** (see **Knowledge Management And Reuse** below). In other words, a firm will have to more or less innovate constantly, and do so within increasingly short spans of time (see **Knowledge Management And The Cash Flow Cycle** below). The key strategic question is not *will we innovate?* but *will we innovate fast enough, often enough and efficiently?* Because knowledge management systems make visible and measurable what is otherwise latent and presumably inefficient, and identify where in fact the knowledge on which innovation is based

comes from inside a firm, and how that knowledge most quickly gets into the hands of the marketplace, they are the key to a program of ceaseless innovation in products, services and processes.

Knowledge Management And Exnovation

By its very nature, knowledge is disruptive and destabilizing. Although the classical definition of the firm is essentially conservative -- the firm is inherently inertial and therefore good -- more modern critiques of the firm as a device for multiplying capital have focused on the notion that firms ought to be courting their own radical transformation, rather than continuing to do what they have always done in the way that they have always done it. BPR is nothing more than a particular view of how exnovation -- the removal of existing structure in favor of other more appropriate structure-- ought to take place. Exnovation, it turns out, is as important to a firm's health as innovation -- and perhaps is a necessary precursor to innovation.

But BPR -- as emblematic of most thinking about exnovation in firms -- says very little about what it is that triggers the drive to exnovation. Again, it is knowledge - - either about an existing structural deficiency or a "better way" -- that triggers exnovation activities in a firm. And, once again, because knowledge management systems make visible and measurable what is otherwise latent and presumably inefficient, and identify where in fact the knowledge on which exnovation is based comes from inside a firm, and how that knowledge most quickly gets into the hands of the people in the firm who execute the exnovation, knowledge management systems are the key to successful exnovation, and perhaps more importantly, to the prevention of exnovations that the firm finds strategically or tactically undesirable.

Knowledge Management And The Cash Flow Cycle

Time-to-market, along with time-to-decide, are perhaps the two most critical cycle time metrics in business today. As time-based competition becomes the norm, particularly where the development and introduction of new products and services are concerned, firms intent on staying in business are going to have to devote increasing amounts of management attention to the concept-to-cash-flow cycle time: the amount of market time required to move from identification of market opportunity through development into delivery and positive cash flow.

For the most part, previous management disciplines have rationalized and routinized the back end of this process: product and service development and delivery models are widely-documented, generally efficient and available.¹³

¹³ And subject to the same disjunctive effects of technology or process innovation, making dinosaurs of efficiency over night. Consider for example the effects of the World Wide Web on the delivery processes for for-fee data.

However, the front end of the process -- market identification, opportunity characterization, and product and service design -- are still knowledge-intensive black arts that appear, from all industry studies available, to be consuming an increasingly large portion of the total concept-to-cash-flow cycle time.

Knowledge management systems applied particularly to the front end of the cycle, focusing on extracting good methods and practices in this area, and welding a well-defined design process to already-efficient development and delivery processes, are likely to produce significant cycle time reductions and increased likelihood of significant positive cash flow every time (fewer failed or misdirected product or service programs).

Knowledge Management And Cultural Hygiene

The problems of cultural hygiene -- of developing and maintaining a healthy corporate culture -- are becoming more intractable as we are forced to deal with a firm's human assets as individuals, as employees and as members of the fluid constituencies that make up a firm at a time when

- those individuals are more vocationally mobile than ever
- basic assumptions about "what employees want" are being overturned and plowed under
- the constituencies employees build are coming 'out of the closet' and into open contestation for the firm's destiny
- the essence of the firm's contract with its employees -- job security in exchange for labor -- is increasingly suspect by employers and employees alike.

In this environment, so-called employees holding so-called "thinking roles" -- knowledge workers whose primary role within the firm is to construct, embody and disseminate knowledge -- become problematic for the firm in two ways:

- the demands of a knowledge worker, in terms of the terms and conditions of her employment 'contract' with the firm, are increasingly idiosyncratic, and increasingly expensive for the firm, leading to a significant net escalation in the burden rate associated with thinking-role headcount
- the contributions of a knowledge worker to the firm can be interrupted without notice, and, more disconcertingly, contributed to a competitor's firm immediately thereafter.

In a firm without formal knowledge management systems, it is nigh unto impossible for human resources personnel and general management to

- assess the demands of this knowledge worker relative to her contribution to the firm's knowledge bases -- is the employee in question worth her cost?
- form an adequate strategic assessment of the damage done to the firm by the loss of the employee in question
- systematically require high-value knowledge workers to embody their knowledge within firm products, services and practices to prevent the loss of intellectual capital when the worker in question leaves the firm
- understand the knowledge worker in the context of her working network, so that promotions and assignment changes do not damage productive knowledge communities within the firm.

All of these necessary cultural hygiene tasks are enabled by knowledge management systems.

Knowledge Management And The Competition For Human Assets

Charles Handy and others have remarked on the fact that, worldwide, we appear to be producing fewer knowledge workers, not more -- that, in fact, the vocational, specialist-oriented curricula that colleges and universities implemented in the 1980s in response to commercial requirements made public in the 1970s are not only not producing the workers needed now, but are producing workers singularly ill-suited for the flexible, 'ceaseless learning' environments of the 1990s.

This suggests that it is increasingly important for firms who expect to capture an unfair share of the knowledge workers in their respective labor pools to offer a working environment amenable to those prospective employees. A firm that does not practice knowledge management (and compensate knowledge workers based on their contribution to the firm's knowledge) will be at a significant competitive disadvantage in such a market.

Knowledge Management And Reuse

Reuse is the holy grail of knowledge-intensive commerce. The mass customization model is premised, fundamentally, on the notion that the only viable business models are those that assume both standard product and service components and "markets of one" -- that reach extremely high levels of component reuse.

Particularly in service markets, much of the “customization” is the adaptation of the “wet” components of the product or service -- the knowledge -- to a particular customer, on top of more or less standard product or service building blocks.

Firms without knowledge management systems will be effectively unable to achieve the reuse levels required by the business model implicit in the markets they enter, and will lose market share to those firms who do practice knowledge management.

Knowledge Management And Corporate Memory

Commercial firms, particularly US firms, have notoriously poor memories. The corporate past is lost almost at the very minute it becomes “the past” -- we cannot remember what we did, the circumstances in which we made decisions, the rationale for those decisions, who said what to whom and what they meant by it. This is exacerbated by our collective tendency to view discontinuity, in the form of mistakes, errors in judgment, or misperceptions as evidence of “bad management” and our resulting predisposition to “revisionist history” and the rewriting of the corporate past to erase all discontinuity and replace it with a smooth “evolutionary view” of decision-making and policy-setting that, while it creates great cognitive dissonance within the rank-and-file, makes us feel better about ourselves when next we speak to the board of directors or shareholders.

The problem with this forgetfulness is that we are unable to learn from the past. Forward thinking firms are, paradoxically, simultaneously embracing learning organization precepts like “fail early and often” and at the same time erasing the traces of their failures through forgetfulness or revisionist myth-making, depriving themselves of one half of the knowledge creation machine: experience. It is truer today than ever that it is easier to know what **not** to do and how **not** to do it, based on past failures, than it is to know what to do and how to do it based either on experience or more conceptual knowledge bases.

Knowledge management systems intervene in this dysfunction by capturing rationale: the context in which decisions are taken, the parties to the decisions in question, and the heuristics used to make the decision. The technology infrastructure knowledge management systems drag with them record history, aid our collective memory and prevent us from engaging in revisionist history.

Bibliography

Chris Argyris. *Knowledge for Action: A Guide To Overcoming Barriers To Organizational Change*. San Francisco: Jossey-Bass, 1993. ISBN: 1-55542-519-4.

Manuel Castells. *The Informational City: Information Technology, Economic Restructuring and the Urban-Regional Process*. London: Basil Blackwell, 1989. ISBN: 0-631-17937-2.

Peter Clark and Neil Staunton. *Innovation in Technology and Organization*. London: Routledge: 1989. ISBN:0-415-00422-5. I am deeply indebted to Clark and Staunton for a number of the fundamental models in this essay.

Peter F. Drucker. *Post-Capitalist Society*. New York: HarperCollins, 1993. ISBN: 0-88730-620-9.

Peter F. Drucker. *Technology, Management and Society*. New York: Harper and Row, 1970.

Michel Foucault. *The Archaeology of Knowledge & The Discourse On Language*. New York: Pantheon, 1972. ISBN: 0-394-71106-8.

Steven L. Goldman, Roger N. Nagel and Kenneth Preiss. *Agile Competitors and Virtual Organizations: Strategies for Enriching The Customer*. New York: Van Nostrand Reinhold, 1994.

Ikujiro Nonaka and Hirotaka Takeuchi. *The Knowledge-Creating Company: How Japanese Companies Create The Dynamics Of Innovation*. Oxford: Oxford University Press, 1995. ISBN: 0-19-509269-4.

Michael Polanyi. *The Tacit Dimension*. New York: Doubleday: 1966.

Michael Polanyi. *Personal Knowledge: Towards a Post-Critical Philosophy*. Chicago: The University of Chicago Press, 1962. ISBN: 0-226-67287-5.

Peter Senge. *The Fifth Discipline: The Art and Practice of the Learning Organization*. New York: Doubleday: 1990. ISBN: 0-385-26095-4.

Solveig Wikstrom and Richard Normann. *Knowledge and Value: A New Perspective on Corporate Transformation*. London: Routledge: 1994. ISBN: 0-415-09817-3.

Charles D. Winslow and William L. Bramer. *FutureWork: Putting Knowledge To Work in the Knowledge Economy*. New York: The Free Press: 1994. ISBN: 0-02-935415-3.