

The Knowledge Value Chain (KVC): How to Fix It When It Breaks

Tim Powell
Managing Director, TW Powell Co.

Presented at *KnowledgeNets 2001*, New York City, May 2001.
Published in M.E. Williams (ed.), *Proceedings of the 22nd National Online Meeting*.
Medford, NJ: Information Today, Inc.

Keywords

Knowledge Management, Value Chain, Knowledge Development, Knowledge Application

Abstract

This paper applies the “value chain” model developed by Michael Porter and others to knowledge-based processes. Parallels are drawn between knowledge development/application and other value-added processes such as manufacturing. As in manufacturing, the Knowledge Value Chain is described as a series of States (stages of processing), with a series of Action Steps required to move from each stage of processing to the next. An example involving the highest levels of strategy in a major US firm is described. Various implications of the model are drawn, including those that use the model to diagnose common problems in the development and application of knowledge.

Background

I had recently joined the Board of the Society of Competitive Intelligence Professionals (SCIP) when I was approached by a new member, who was also new to the field of competitive intelligence (CI). She asked me, “What do CI people actually do?” I began explaining to her the value-adding transformation of data to information to knowledge, and finally the communication to someone with the authority and power to do something about it (a “decision-maker”). I advised her that thinking of it like any other manufacturing process would help de-mystify the process, and give her a grounding in the “real world”. Then I began to diagram, on the back of an envelope, the beginnings of what later became the fuller version of the Knowledge Value Chain presented here.

Introduction

Virtually all of us reading these words are “knowledge workers” — in doing our jobs, we read articles and journals, search databases, talk to people, think about and summarize what we’ve found, and communicate our findings to decision-makers. But the *process* by which this occurs is not well-defined or understood, either by cognitive scientists or by business process experts. Yet, the knowledge-creating process stands at the heart of modern organizational competitiveness.

We must understand the knowledge process before we can attempt to improve on it. We must peer into the “black box” of knowledge work and attempt to see what is inside, and how it works. Why? Simply put, the productivity of the knowledge worker is seen by experts¹ as the single most important factor in the competitiveness of the modern organization. And productivity is best enhanced when work processes are broken into their basic elements.²

We must find some way to turn *information* into *business value*. Stated that way, this transformation sounds magical, like the transformation of lead into gold that the medieval alchemists struggled to achieve. Let’s not forget that their struggle, though ultimately futile, did produce the body of knowledge that eventually became modern chemistry. That’s because they began to document, test, and refine each micro-process that led in the direction of their final goal, and to assemble those “micro-learnings” into a larger body of knowledge.

That is what we will begin to do here. I should point out that this thinking builds on the work of others. And I hope it will not be the last word on this subject, as any model of this nature needs challenging and refining before it becomes truly useful. As such, I claim no proprietary right to this term Knowledge Value Chain, and in fact hope it will be widely adopted.

Definitions

In defining the Knowledge Value Chain, we need to define each term that goes into it. Let’s start in the middle, and work our way out.

Creating *value* is ultimately what organizations are all about. In business, we measure value by a number of metrics, each of which is valid when used in the right context. These include:

- ❑ financial measures like revenues and profits;
- ❑ stock price (known in polite circles as “shareholder value”);
- ❑ innovation, for example the number of new patents granted;

- other more complex and sophisticated measures such as the Balanced Scorecard developed by Nolan and Norton, and Economic Value Added® (EVA) developed by Stern and Stewart. These include a mix of measures like customer satisfaction and production quality.

Non-business organizations have similar measures of effectiveness and efficiency.




Figure 1: A Simple Value Chain

The *value chain* model was popularized by Michael Porter in his 1985 book *Competitive Advantage*³. We'll define it here as *a series of related activities that together produce end-user benefits*.⁴

A simple example of a value chain

is manufacturing. Here there is a defined series of steps in which raw materials are transformed into work-in-process at various stages of completion, and then into finished goods ready to be packaged and shipped. Each stage of processing adds both costs and value (defined here as the purchase price paid by the end user.) Each step should add more value than cost – and if not, that step needs to be re-engineered in some way.

An illustration from everyday life is painting your apartment or house before you try to sell it. People usually do this, as typically the cost of doing so is less than the value added – even if the new owner intends to repaint before moving in.

The *knowledge value chain* is simply the application of the value chain model to knowledge development and usage. In short, it is a process model of how data becomes intelligence, and eventually becomes part of a business result or benefit. The balance of this paper will define this KVC process.

The Knowledge Value Chain

The knowledge value chain (KVC) comprises two major sets of activities, *knowledge acquisition* and *knowledge application*. This reflects the division of labor in knowledge work that has evolved in large, complex organizations. Here, *knowledge workers* are primarily tasked with knowledge acquisition and development, and *decision-makers* apply the resulting knowledge to make better business decisions, plan and execute actions, and thereby achieve business results.

The KVC model contains an alternating series of “states” and “actions”. Each state is a *stage of processing*, a work-in-process like a partially-built car would be in an auto factory. Each action is the *transformation* needed to move to the next stage of processing.

There are a total of eight process states and seven key transformations or action steps. Each action step has as its “input” the state preceding it, and as its “output” the state following it. Each of the seven steps is an opportunity to add value – or to fail to do so,

or even destroy value, if it is omitted or executed poorly. Each step is thus a potential *fail point*, at which we need to be asking whether we are controlling quality.

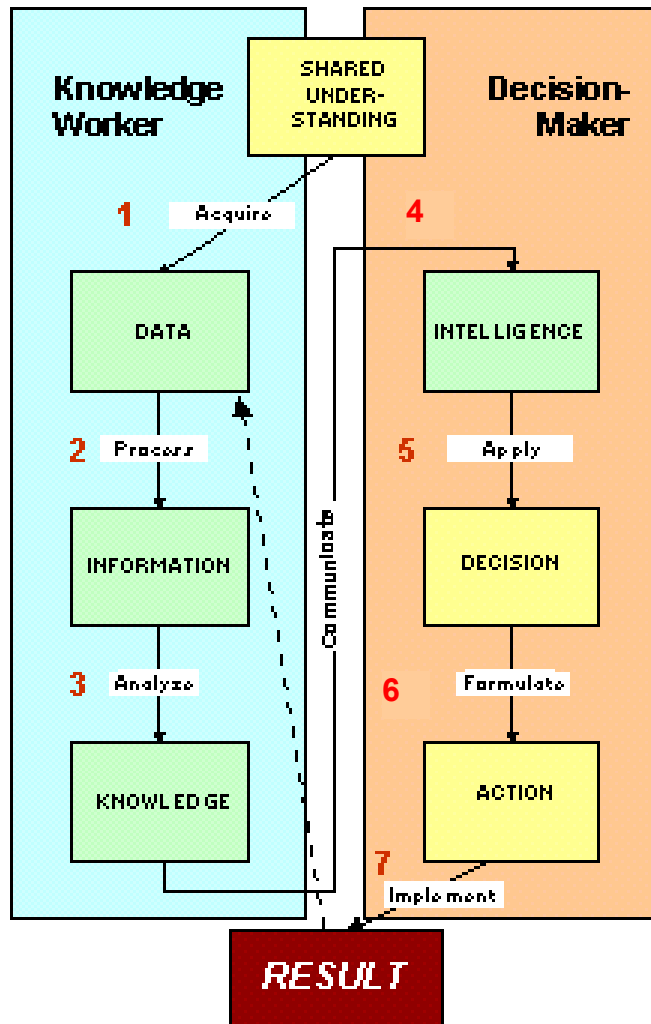


Figure 2: The Knowledge Value Chain (KVC)

The KVC model starts with a *shared understanding* between the knowledge worker and decision-maker. The quality of this step is essential to the whole rest of the process. If this part is flawed (for example, rushed--or skipped entirely), the remainder of the process cannot achieve the level of value-enhancement possible.

To use our car manufacturing analogy, in some cases great expense and time are spent designing and manufacturing cars that ultimately fail in the marketplace because the "decisions maker" (that is, the consumer) does not want them. By contrast, the Ford Taurus is legendary in automobile history for two reasons: (1) it was one of the best-selling models year after year, and (2) massive amounts of consumer research went into defining the car's features before it was even designed.

In the world of business research, an example of a lack of shared understanding is when you spend a lot of time researching a problem that turns out not to warrant the urgency you gave it.

Both examples derive from a failure to achieve shared understanding at the outset. This understanding should include the basic project parameters, such as:

- ❑ the time frame,
- ❑ the budget (if any),
- ❑ the use that is to be made of the information, and
- ❑ any sensitivities that may exist around the project.

Action Step 1. Acquire Data

Then, and only then, can the knowledge worker go off and begin to *acquire* data. Data are the equivalent of “fact and figures”, the points around which we will eventually build our analysis. (When we say acquire, we are speaking of any method of obtaining data, including but not limited to its purchase for a fee.) This can involve speaking to people on the phone, going online, looking at printed sources, and so on. (Exactly where and how to acquire data is a long discussion – refer to Analyzing the Competition, for example.)

Given the proliferation of data that exists, as often as not this step involves selecting and filtering a bounded set of data from a much larger data universe.

Having acquired the data, then the analysis can begin...almost.

Action Step 2. Process Data

But first, in my experience there is usually an interim step that involves *processing* or organizing the data – preparing it so that it can be analyzed properly. This step frequently involves entering data into a computer. For example, handwritten notes from an audiotape of a phone conversation might need to be transcribed into a word processor before they can be analyzed. Numeric data points can be entered into a spreadsheet, and trends lines and charts drawn. Publications in a foreign language must be translated into the analyst’s native tongue. Facts need to be checked and validated. “Outliers” – data points that lie far outside the general trend of the data – need to be examined, and decisions made as to whether they are valid and should be included, or whether they are spurious and should be dropped.

Action Step 3. Analyze Information

The processed data is called *information*, and only now it is ready to be analyzed. During the *analysis* phase, we transform information into *knowledge* in the following ways:

- ❑ by giving it *context* – what are the circumstances surrounding it?
- ❑ by assessing its *relevance* to our organization – what does it mean *for us*?
- ❑ by drawing *implications* for action – what do we recommend be done about it?
- ❑ by outlining *options* and alternatives, and
- ❑ by *reducing* the amount of information to an amount able to be understood and acted on by the decision-maker. In this fast-paced world, *attention* and *focus* are two very precious commodities in the world of even the highest-bandwidth decision-maker, and it is the analyst’s job to produce both.

Action Step 4. Communicate Knowledge

In this most crucial step, knowledge that heretofore has been in the mind of the knowledge worker is transmitted to a decision maker (or, more typically, decision-making group). This typically happens by a combination of means including e-mails, formal written reports, slide presentations, and personal briefings. In effect, this step replicates the knowledge-building process in other individuals grappling with some business problem. Here, though, the process must be accelerated and concentrated.

A successful communication gives the right amount of material, summarized and organized well, and with analysis, conclusions, recommendations, data sources, and any qualifications needed on the quality of the data or likelihood of outcomes predicted (“If presents trends continue...”, “Experts predict that...”, etc.) A successful communication produces *insight* in the decision-makers, and *credibility* for the knowledge worker.

This transferred knowledge is now *intelligence*, defined as *knowledge in the hands of someone with the capability of acting on it*. Once this hand-off to a decision-maker is made, too often a knowledge worker will consider the job done – and in an active sense this is true. But in the sense of value-creation, the process to this point has not created business value, since nothing has yet been accomplished with the information.

Knowledge workers who seek to improve their effectiveness need to know what happens after this knowledge transfer. Only by knowing what is done with the knowledge that we give to decision-makers will we be able to improve our work product and processes. Just as market researchers find out what happens to products once they go into usage, the knowledge worker needs to understand the process after he/she is directly involved.

Action Step 5. Apply Intelligence

Decision-makers do just that, take intelligence, and use it to make decisions for the organization, most often in a group setting. These decisions typically involve resource allocations – where to invest, where to pull back, etc.

A key factor in any decision is when to make it. Some decisions are *time-driven*, for example budget decisions. Others are event-driven, for example, what to do when a key executive resigns. Other decisions are neither time- nor event-driven, but are driven by change in some variable, for example a falloff in sales. Absent a pre-defined threshold condition (for example, review sales if they fall by 5% or more), these latter decisions are problematic precisely because it is not clear when they need to be made.

The process we have described so far is a formal decision-making process. Study after study has shown that, while such formal processes are often involved in any major decision, there are also typically informal inputs. These can include the decision-

maker's former experiences, informal sources such as a spouse or neighbor, intuition, and a whole host of "externalities" such as political factors.

Significant barriers can come between the presentation of good intelligence and its application. Failures to apply intelligence can include ego ("Not invented here"), group pressure ("They all think it's the way to go, maybe they're right"), and what they psychologists call *cognitive dissonance* ("Don't confuse me with the facts").

Action Step 6. Formulate Action

We are getting closer to the end results, but are not there yet. We need to plan and execute some actions that embody the decision. This step is a major source of "disconnects" in the KVC. For example, in the US our lawmakers and the courts stop at decision-making, and leave it to an army of federal agencies to try to implement their decisions. Even when a good decision is made, it can be rendered totally ineffective by poor actions planned to implement it.

Action Step 7. Implement Actions

Now the actual implementation of the decision and resulting action plan begins. Resources must be committed, goals and timetables must be set, and some agency for monitoring the results must be devised. Now, and only now, are we in a position to gain business results from our original knowledge work.

This is also a key chance to gain some new data. Note the dotted-line feedback loop in the diagram above. As implementation goes forth, there is an opportunity to gather data of a typically higher quality than in preceding stages. This feedback needs to be put into the cycle at the top, and analyzed quickly in order to provide a "course correction" to the whole process.

For example, the credit card company Capital One is known to have an extremely effective marketing operation built on direct mail. While all the card companies do a lot of direct mail, Capital One has gotten really good at doing many focused tests ("Would people in zip code 10014 who recently bought new computers be interested in our card?") The results of these tests are fed back in order to continually refine the whole process.

A KVC Illustration

We can use the mnemonic "SDIKI-DAR" (rhymes with "sticky tar") to guide us through the KVC process. As an illustration, let's apply the KVC model to a company's decision to acquire another. *Shared understanding* could include a basic definition of the characteristics of a target company of interest—industry, revenues, profitability, market

share, distribution, etc. *Data* might include the stock price and trading volume of a potential target. *Information* (processed data) would include trends lines, arrays, and graphs built out of this data. *Knowledge* could include, for example, correlations between the stock price and volume trends with news about the company, and with conditions in the market in general.

Intelligence would result from the analyst's bringing all of this data to the attention of a decision-maker. A *decision* could result to make an offer to acquire the target company. An *action* plan would be formulated to define the terms, process, and support players for the acquisition. The ultimate *result* would include the impact of the acquisition on earnings per share and stock price, and the acquiring company's market share in the affected market.

Typical KVC Disorders

As mentioned earlier, the point of understanding the KVC process in detail is to be able to fine-tune it and fix it when broken. I've recently had the opportunity to study competitive knowledge development in a wide range of organizations and industries. In doing so, I've been struck by the fact that knowledge workers across industries and geographies have an amazingly similar set of challenges and problems. I believe that many of these can be readily understood and corrected by using KVC analysis. Several examples follow.

PROBLEM: *"We don't know what management wants."* This is a common complaint of knowledge workers. Clearly this is an "Action Step 0" problem (that is, pre-shared understanding) because there has been an inadequate dialogue between knowledge worker and decision-maker. As a result, the data to be acquired can't be defined properly, and the whole chain spins off its axle. The reader may chuckle at the words "inadequate dialogue" – indeed, *dialogue* glorifies what is often an eleventh-hour frantic phone call for help.

SOLUTION: *Make sure a true exchange happens before data collection starts. Define the project parameters listed above. Use a checklist if you're afraid you'll forget something. As a rule, project misunderstandings are best had up front, so they can be corrected early.*

PROBLEM: *"We have piles of data – but never enough time to analyze it."* This is an analysis (Action Step 3) problem, and commonly results from the rush to complete projects. If you fail to process and/or analyze data, you will be tempted to provide unprocessed, un-analyzed data or information to a decision-maker. This reduces the value of the communication, and if done repeatedly calls your credibility as a knowledge worker into question. You must add value through analysis – though it is always good to separately identify "the facts" from your analysis of them.

SOLUTION: *In planning the project, be sure to allow adequate time for processing and analysis. Back-time from the "deliverable" deadline, including the final production, and the analysis to get to the results. Always allow some "slip time", as Murphy's Law is alive and well. ("If something can go wrong, it will. If something cannot possibly go wrong, it still will.")*

PROBLEM: “We do great analysis – but management doesn’t use it.” This communication (Action Step 4) problem results when credibility has not yet been established, or the relationship with the decision-maker has been damaged in some way. Creating a positive relationship and ongoing dialogue with decision-makers is one of the most difficult and important challenges in any knowledge-based process. Knowledge workers and decision-makers co-exist across a cultural gap that often impedes their communications. Almost everything about them is different. They typically are trained differently, paid differently, dress differently, live and work in different places, and so on. The need for and difficulty of bridging that cultural gap, and to gain trust and credibility, is often underestimated.

SOLUTION: Do some “market research” with your decision-maker clients. Determine what they find most effective in terms of content, format, and medium of communication. Follow up on projects – what was most useful, what was least useful, what needs to be improved.

PROBLEM: “We do great analysis – but management keeps asking for more.” This is an Action Step 5 problem (decision), in that decision-makers may feel that they need a greater level of certainty before making a decision. As a rule, the greater the level of resources to be allocated based on the decision, the greater the level of certainty that decision-makers will seek. But, paradoxically, a good data collection and analysis effort often raises as many questions as it answers. This leads to the apparent need for more information. Typically, though, the costs and time to produce new analysis (and especially new data) are underestimated, and can impede the smooth flow of decision-making. This can lead to the situation commonly known as “paralysis by analysis”.

SOLUTION: Raise the sensitivity to the trade-offs between greater certainty, and greater time, effort, and budget expended. Keep the flow of events moving forward toward decision-making.

PROBLEM: “Once we go into implementation mode, we stop collecting data.” The feedback loop (post-KVC) is often ignored, thereby cutting off a source of real-world validation and course-correction for the original decision. The problem may be an organizational one – the issue is handed off from a “research and development” team to an “implementation team”. The latter is less familiar with data collection, may not have the resources for it, and generally does not have the cultural bias toward information gathering that the research people do.

SOLUTION: Create an ongoing, active communications linkage between the research and implementation teams. Encourage the feeding back of new data into the KVC.

PROBLEM: “We have many KVCs – but they don’t link with each other.” One key difference between the KVC process and a physical manufacturing process stems from the intangible nature of the KVC “products.” In an auto factory, once a bolt has been used in one car, it obviously can’t be used in other. But data and information are “non-wasting” assets, they don’t get “used up” when used. As such they are infinitely re-usable. Data used in one KVC can be used in another somewhere down the road – and often should be. Too often this doesn’t happen due to the boundaries of time or organizational structure.

I recently conducted a knowledge development workshop with a major pharmaceuticals company. In one of the group exercises, we determined that the competitive product information gathered by scientists at the beginning of the drug development process stayed in the R&D unit. We discovered that many months later, when the drug had completed development, testing, and approval, the same data could be useful to the sales force tasked with moving the product into doctors' offices.

We transferred the same data (comparative drug effectiveness) from the R&D teams to the sales teams. Building this simple linkage saved the re-creation of the information at a later time and different context.

SOLUTION: Create a macro-view of the various knowledge chains in your organization. Look for opportunities to re-deploy data in ways that enhance value.

A Tragic Example

Over time, the misuse of knowledge can erode the fundamental value created by a business. A recent example, tragic in its proportions, comes from the Xerox Corporation.

Until recently, Xerox was considered a paragon of corporate practice in many respects. For example, they pioneered some competitive applications of information—the modern practice of benchmarking was perfected there.

Lately, Xerox stock has recently been trading at a fraction of its recent high, has developed serious credit problems, and has repeatedly been rumored to be going bankrupt.

Xerox's organizational use of knowledge is at the core of their fall from grace. During the 1970s, scientists at their Palo Alto Research Center (PARC) were responsible for a stunning series of key inventions of the digital age: the laser printer (1971); ethernet (1973); the graphic user interface (1975); the personal computer (1977); and the portable computer (1978).⁵ Yet in every single case, they were ultimately unable to use these inventions to create a dominant position in the commercial marketplace. What's worse, in each case another company *was* able to do so. For example, Steve Job's emulation of the Xerox Star computer user interface technology in the Apple Macintosh is now legend. IBM (now Lexmark) and Hewlett-Packard took most of the share of the laser printer market. Ethernet was commercialized by Intel, Digital, and 3Com.

The Xerox problem was a variation on the knowledge worker-decision maker gap described earlier. The scientists at PARC were removed from the company's primary decision-makers both physically and culturally. The PARC scientists had a shirt-sleeve Silicon Valley culture long before it was fashionable in the East, and were driven by creativity—and even fun. The Stamford, CT decision makers worked in a blue-chip, button-down, bottom-line environment. They were skeptical of the new ideas from the

Wild Westerners, and ultimately refused to fund the commercial development of these key technologies.

This series of decisions not to enter new markets, overlaid with drastically changing conditions in Xerox's core copier market, created long-term strategic problems from which the company may never recover.

Key Implications

Now that we understand the KVC process in principle, how can we apply it to solve problems? What lessons and implications can we draw? Some of these are as follows:

- ❑ *Knowledge is a linear process.* In contrast to some “cyclical” models of knowledge development, KVC is a linear model. Like manufacturing, it uses a series of steps to define a final product. The key lesson here is that *you must understand the destination before you start the journey*. And that destination is not simply the approval of your decision-maker client (though that may be a good real-world proxy for the true payoff.) The true payoff is the final business condition that will result from the whole process.
- ❑ *Knowledge is a serial process.* These states and actions are “wired in series”. Picture those holiday lights on strings, where when one goes out, the whole string goes out. Likewise, here, *when one link in the value chain is broken, by definition the chain itself is broken*. A failure at any step is replicated in all later steps. For example, a great analysis of faulty data will produce a misguided conclusion. The old adage “garbage in, garbage out” makes this point very succinctly.
- ❑ *Each step in the KVC is integral and essential.* If you short-circuit the process (which you might be tempted to do, to save time or other resources), you will fail to create value – or possibly even destroy value. The classic example is when a decision is made without sufficient information to support it. The result is more a matter of luck than of strategic management. While it is possible to succeed in the short term under these conditions, over the long haul the odds are stacked heavily against doing it repeatedly.
- ❑ *Knowledge value chains interact.* In the pharmaceuticals case mentioned above, a single data set was employed in two otherwise separate KVCs. You need to map out these various KVC processes, formally if possible, in order to identify and improve these interactions. As an example, when you finish a project for a decision-maker, it is helpful to routinely inquire as to who else would find the information useful (with, of course, your client's advice and consent.)

- *Knowledge creation is like manufacturing.* I state this here as a conclusion, though admittedly it was more accurately a hypothesis going into this series of propositions. However, to the extent that it is true, we can then use the many insights into manufacturing that have been gathered by students of that discipline.

Manufacturing concepts that we can “borrow” in the study of knowledge-based processes include:

Incremental value added. Each of the seven transformation steps can be deconstructed, and some assessment made of its incremental cost and value-added. Though it may not be possible to do this on a strictly quantitative basis, at least at first, the exercise will still yield important insights.

Productivity. The productivity of knowledge-based processes is a primary aspect of competitive differentiation. In manufacturing, it can be measured the by ratio of input to output (for example, person-hours to produce an automobile.) Again, the measure for knowledge-based processes may be a lot more “fuzzy” at this stage than it is for manufacturing, but the model may prove as valuable.

Quality assurance. Just as in manufacturing, quality can be built into the process. The alternative can be to find out when a complete study has been done, and resources expended, that they are the right answers — *but to the wrong questions*. One form of quality assurance, for example, can be establishing frequent touch-points with the decision-maker on the project status (from both sides.) This will avoid any “process surprises.”

Load balancing. Too often a knowledge process is off-balance in terms of resource assignments. A common example is that too much time is devoted to data collection, and too little to analysis and communications. Load balancing enables you to identify and clear the bottlenecks that may be dragging down the whole process, and to re-allocate resources in order to clear those bottlenecks.

Just-in-time (JIT). JIT manufacturing came into vogue because associated with each raw material and purchased component, there is a cost of inventory (typically, the financing cost of purchasing the items, plus the cost of the physical warehouse space in which it is kept.) For information, the inventory cost is typically not the “warehousing” cost — though there may be some cost associated with physical document storage and magnetic media such as hard drives.

The main cost associated with information is its *obsolescence cost* — its inexorable tendency to become less representative of the “real world”, hence less valuable, over time. Any piece of produced information begins to “decay” as soon as it is produced. Some information decays rapidly (“Who wants yesterdays’ paper?”), some less rapidly.

As a result, it is useful to consider production of knowledge at the last possible moment of knowledge, to minimize the risk of obsolescence. (Come to think of it, maybe that's the reason for all those last-minute information requests!)

Inventory. The view of information as a highly perishable asset, like fresh fish, raises other considerations. Issues include: which information should we keep "in inventory"; which should we produce on demand; what is the current value of what we have in inventory; etc. The nature of the inventory can vary, for example, a current list of print sources may prove nearly as valuable as having the sources themselves--assuming those sources are readily retrievable when needed.

Auditing. Physical inventories of raw materials, work-in-process, and finished goods are subjected to periodic structured verification processes called *audits*—and our knowledge inventory should be, too. In this way you can determine what you have, what condition it is in, what you have too much of, what gaps there are, and so on.

Though information is different from other assets in key respects (primarily, its intangibility and re-usability), it is nonetheless useful to consider those ways in which it is *similar* to other assets. In this way, many pre-existing management models, some of which are listed above, can be brought to bear on the strategic problem of knowledge productivity.

Conclusion

"The purpose of information is not knowledge. It is being able to take the right action."⁶ Too often knowledge workers limit themselves to the production and distribution of knowledge, without considering how it creates value for the organization. This key omission threatens the value of knowledge work—and the careers who perpetuate this way of thinking.

We have presented the Knowledge Value Chain model in the hopes that it will spur enhancements in the productivity of knowledge processes—the key determinant of organizational competitiveness. We do not intend it as a theoretical model, but rather for use on the "front lines" of knowledge about knowledge—in diagnosing and fixing problems.

Notes

¹ For example, Peter Drucker: “Fifty years from now – if not much sooner – the leadership in the world economy will have moved to the countries and to the industries that have most systematically and most successfully raised knowledge-worker productivity.” In *Management Challenges for the 21st Century*, Harper Business, 1999, page 158. I believe the same statement will be applicable to *companies* and *individuals* as well, and that the time frame will be considerably shorter than 50 years.

² Frederick Taylor did this for industrial work. See Andrea Gabor, *The Capitalist Philosophers*, Times Business, 2000.

³ Michael Porter, *Competitive Advantage*, Free Press, 1985.

⁴ Note that this is similar to the definition of a *business process*, and indeed the concepts are closely related.

⁵ Michael Hiltzik, *Dealers of Lightning*, Harper Business, 1999.

⁶ Drucker, op. cit., page 130.