



Meeting the Challenge of the Shift to Knowledge Work

Fred Nickols

Foreword

I first wrote this paper in 1983. I came across it recently while digging through some old files and, after reading it, I decided the issues and ideas are still relevant and that it was worth posting to my articles web site. That it's still relevant after 20 years probably says something about the rate of progress that's been made on the issues. Indeed, a study cited in the December 2003 issue of *Training* magazine characterizes the state of performance among knowledge workers as "sunken."

This is the first time the original paper has been available in its entirety. A few of the ideas in this paper appeared in the October 1983 issue of *Performance & Instruction*, in an article bearing the title, "Half A Needs Assessment: *What Is* in the World of Work and Working." That article, minus its needs assessment introduction and with a few minor changes, was reprinted as a "classic in knowledge management" in the 2000 edition of the Butterworth-Heinemann *Knowledge Management Yearbook*.

In the course of keying the original paper into my word processor, I fixed a few typographical and grammatical errors, corrected some awkward wording, inserted a few footnotes and devised a new opening but have otherwise remained faithful to the original. The original paper begins below.

The Shift to Knowledge Work

That there has been a shift to knowledge work is easily demonstrated. Peter Drucker, noted professor, economist and consultant, has chronicled it over the years (see the box above)¹ and its size and swiftness are reflected in statistics drawn from the Bureau of the Census and the Department of Labor.

In 1900, almost three-fourths (73%) of the work force consisted of people engaged in manual work (i.e., farm or factory work). Roughly one-quarter (27%) was engaged in knowledge work (i.e., office or service work). By 1920, the shift was underway: two-thirds of the work force consisted of people engaged in manual work and one-third was

Peter Drucker's Chronicle of The Shift to Knowledge Work

"In the United States...the class of employees that has been growing most rapidly in numbers and proportion is that of skilled and trained people."

The Practice of Management - 1954

"Even the small business today consists increasingly of people who apply knowledge, rather than manual skill and muscle, to work."

Managing for Results - 1964

"Finally, these new [knowledge] industries differ from the traditional 'modern' industry in that they will employ predominantly knowledge workers rather than manual workers."

The Age of Discontinuity - 1968

"...the center of gravity of the work force is shifting from the manual worker to the knowledge worker."

Management - 1973

"...the center of gravity among 'employees' has sharply shifted to the educated, employed middle class, that is, to people who see themselves as 'technical' and increasingly as 'professional.'"

Managing in Turbulent Times - 1980

¹ Drucker, of course, continued chronicling and commenting on the shift to knowledge work well past 1980. An extended version of this chronicle can be found on my articles web site at <http://home.att.net/~nickols/chronicle.htm>.

engaged in knowledge work. By 1980, things were the other way around: two-thirds of the work force was engaged in knowledge work and only one-third was engaged in manual work. In the brief space of 60 years, work changed from a materials-base to an information-base, changing forever the nature of organization and the role of the worker in it.

The shift to knowledge work is even more dramatic in our “high-tech” or knowledge industries. In testimony before a Senate subcommittee looking into employment and productivity, David Dallob, vice president of Sperry Corporation, estimated the percentage of knowledge workers in the computer industry at 72%. In the consulting industry, the percentage of knowledge workers is probably close to 100%.

The supposed stronghold of manual work – manufacturing – is no exception. In testimony before the same Senate subcommittee, David Birch of MIT indicated that the percentage of the work force still engaged in actual manufacturing operations is “about 13%.”

Even the military has been affected. In 1956, when I reported aboard my first destroyer, knowledge workers (i.e., clerical, technical and officer personnel) made up about half the crew, reflecting the makeup of the national work force at that time. But, by 1974, when I reported aboard my last destroyer, knowledge workers made up two-thirds or more of the crew.

We are indeed a nation of knowledge workers.

The Challenge of the Shift to Knowledge Work

The challenge presented by the shift to knowledge work is nicely illustrated by two sea stories from my Navy days.

Sea Story # 1: “Aye-Aye, Sir” – The Limits of Authority

The year was 1957. The ship was the *USS Gregory (DD-802)*, an old WW II Fletcher-class, 2100-ton destroyer. The place was the 01-level, back aft by Mount 53. (That’s one deck up from the main deck, and back by the third of five five-inch gun mounts.) We were in Subic Bay, in the Philippines, taking a break from our assignment of patrolling the Formosa Straits.

Tommy Lee Crabtree, a Gunner’s Mate second class (GM2), was working on Mount 53, trying to repair an as yet unidentified malfunction. I was new on board and I was working on Tommy Lee – trying to convince him to invite me to join the armory coffee mess. The armory coffee mess was the most prestigious coffee mess on board my destroyer and I badly wanted an invitation to join. The invitation had to come from Tommy Lee; he was the Gunner’s Mate in charge of the armory. Short-term, my hopes weren’t high but I was prepared to hang in there for the long haul.

Tommy Lee and I were taking a break, hunkered down on our haunches next to Mount 53, sipping coffee and chatting in a way calculated to help him take my measure, when we spotted our division officer approaching.

Our division officer was a Lieutenant - Junior Grade (Ltjg). His last name was Wilson. A bit of a martinet, he had been nicknamed “Whip,” an appellation borrowed from a star of western movies of the 1940s.²

“What are you two doing?” he demanded.

“Drinkin’ coffee and shootin’ the breeze,” replied Tommy Lee.

“What are *you* doing here?” Whip asked of me.

As a Fire Control Technician – an “FT” – someone whose own systems controlled the gun mounts during combat, my work required close coordination with the Gunners Mates so I had a convenient and true cover story. Standing up, I said, “I came down to find out when Tommy Lee thinks we’ll be able to include Mount 53 in the daily workouts and if he thinks we’ll have to realign it with the rest of the gun battery.”

“Well,” demanded Mr. Wilson, turning to Tommy Lee who was still squatting, “when *will* it be fixed?”

“I dunno. I’m workin’ on it. Probably some time today.”

“That’s not good enough!,” snapped Mr. Wilson. “Get off your duff and get back to work! I want that gun mount back in working order A.S.A.P.!”

Tommy Lee gave Mr. Wilson a strange look and then, rising slowly to his feet, Tommy Lee grinned a wicked grin and asked, “Are you *ordering* me to fix this here gun mount, Mr. Wilson?”

“Yes, I am,” said Mr. Wilson.

Shifting his coffee cup to his left hand, Tommy Lee saluted smartly, and said, “Aye-aye, sir. What would you like me to do first?”

The reactions played across Mr. Wilson’s face like moving scenery: first puzzlement, then comprehension, followed in quick order by surprise, shock, humiliation and, finally, red-faced, apoplectic anger.

Tommy Lee did what all those who must submit to authority have been doing for thousands of years, he submitted. He went passive. He told Mr. Wilson to tell him what to do and he would do it. The problem for Mr. Wilson was that he couldn’t issue the necessary orders. He didn’t know how to fix that gun mount and Tommy Lee did.

Furious, Mr. Wilson glared at Tommy Lee, then turned and stomped off without a word.

² Actually, I’ve disguised our division officer’s identity to spare him any embarrassment from my telling of this sea story.

Sea Story #2: Figuring Things Out

It is now a decade later. The place is the Gulf of Tonkin, off the coast of Viet Nam on a summer day in 1967. My ship, the USS Waddell (DDG-24), is providing gunfire support to troops ashore. I'm in the plotting room, in charge of the ship's 5"/54 guns, its main battery. We are in the midst of a mission.

The Gunnery Liaison Officer (GLO), up in the Combat Information Center (CIC), informs me that the spotter, in a light plane a couple of miles inland, has just made an odd request and the GLO wants to know if we can respond. The spotter's request is for us to "throw a couple of rounds as far down the gun target line as we can throw 'em." It seemed the spotter had a target that he suspected was out of range but he wanted to make certain. So, he asked for two rounds at maximum range along the same line we had just been shooting.

Maximum range, in a vacuum, is achieved at a gun elevation angle of 45 degrees. I couldn't do anything about not being in a vacuum; the atmospheric conditions were what they were and they had been entered into the computer and compensated for in its calculations of gun orders. The more vexing problem was that, owing to the pitch and roll of the deck on which the guns were mounted, there was no way of manually setting the guns at that angle and keeping them there. But the gulf waters were especially calm that day. As I observed the dials indicating pitch and roll I could that pitch (up and down) was not a factor and roll (side to side) was less than one degree. At 45 degrees of elevation, it would take more than a degree of roll to shift the mean point of impact of any rounds fired by 100 yards. Satisfied that accuracy would be acceptable, I removed the amplifier that controlled gun elevation order module and hand set it to 45 degrees. In short, I "jury-rigged" the computer controlling the gun mounts. I then reported ready to the GLO.

We fired one round then waited. The spotter came back with a request for six rounds of rapid continuous fire; then six more behind that. Suddenly, over the radio circuits, we heard him exclaim, "Hot damn; secondaries." He was referring to secondary explosions caused by our gunfire. It turns out we hit an ammo dump. All with the gun elevation order manually set to 45 degrees.

Later, done for the day, we steamed off to a different assignment (with the amplifier reinserted in the gun fire control system computer, of course). As we steamed off, I wondered just how far those rounds had gone, so I asked the GLO if he could back-track, checking the charts used up in CIC to see where we were at the time, where the coordinates for the target were, and what the distance was between them. In a few minutes, he told me that the distance to the target we fired at and hit was just over 26,000 yards. I was flabbergasted because that was 3,000 yards farther than those particular guns were supposed to be able to shoot and 6,000 yards farther than the gun fire control system was designed to use. The supposed maximum usable gun range of that system was just under 20,000 yards.

I spent the better part of the night poring over the manuals for that system, trying to figure out why we could shoot so much farther than the guns supposed maximum range. Finally, as I sat looking at the range tables for the gun, I gave up. I closed the range table book and, as I looked at the cover, the nature of the problem dawned on me.

The cover of the range table book indicated that the initial velocity (i.e., the speed of a bullet upon leaving the barrel of the gun) was 2,500 feet per second. Well, that particular gun had an initial velocity of 2,650 feet per second. The range tables were wrong. The design of the gunfire control system was flawed. So, I spent the better part of the following day calculating some crude adjustments to the range table values and converting those to a system of elevation spots that could be inserted into the computer to take advantage of the unused capability of our gun system. We then spent several days blowing up things the enemy thought were safely out of reach.

The "extended range procedure," as I called it, was written up and submitted to the admiral in charge of the task group of which my ship was a part. A young lieutenant on the admiral's staff threw it in the trashcan because he couldn't believe that (1) the design of a Navy weapons system was faulty or that (2) a mere chief petty officer had found that flaw and developed a workaround. (The admiral later asked me to overlook the lieutenant's youth and inexperience, which I did.) Eventually, the gunfire control system in question was modified so as to take advantage of that extra 6,000 yards.

The Central Issue

Both of the preceding sea stories illustrate the central issue in the challenge of the shift to knowledge work; namely, namely, that the distribution of knowledge among the members of our modern, information-based organizations and institutions has resulted in an accompanying shift in the locus of control over the activity of working and thus altered the nature of managerial and organizational control over work. This is an entirely new state of affairs in the annals of collective human endeavor. How well we cope with it will determine our future success as individuals, organizations, institutions, and as a nation.

A modern destroyer, like the one off the coast of Viet Nam, is a "high-tech" organization, a marvelous maze of systems, subsystems, components, communications and people. Its work is information-based and its crew consists primarily of knowledge workers. But that same destroyer is embedded in a "low-tech" culture, that is, a management system built for manual workers.

It has been said of the Navy that it was designed by geniuses to be run by idiots. The same can be said of many other organizations and institutions. In one sense, that's the problem: We've got a high-tech work force caught up in a low-tech management culture. Our ships of business and industry were built to be crewed by muscle workers. So was our ship of state. Guess what? We've got a new crew.

The task of management, then, has shifted from one of obtaining, ensuring and enforcing compliance to one of eliciting and capitalizing on the contributions that only the knowledge worker can make. Said a little differently, *the task of management has become one of managing people who must manage themselves.*

The Nature of the Challenge

The shift to knowledge work radically alters the nature of control over the work of our organizations and institutions. Tight, centralized control over people who use their brains instead of their brawn is simply not possible. Knowledge work (e.g., clerical, managerial, technical and professional work) is done using one’s head. This means the activity of working is no longer visible. As a result, the shift to knowledge work snaps the most critical link in the chain-of-command: supervision.

The locus of control over working has shifted in large part from the management to the worker. It seems safe to assume that this shift in control affects the performance and productivity of our organizations and institutions. The achievement of high levels of performance and productivity is imperative. Society depends on its organizations and institutions to accomplish those ends that can be accomplished only through collective human endeavor. Our organizations and our institutions are also the mainstays of our economy and thus our individual prosperity. If they fail, so do we all.

The challenge, then, is to make our nation productive. Because it is now a nation of knowledge workers, the nature of the challenge is one of making knowledge work productive. And, because the knowledge workers are now in control of working, it falls largely to them to make knowledge work productive. They are the ones who must do it.

The issue of accountability is important. Nothing gets done unless someone is assigned or assumes accountability for making it happen. Ownerless problems don’t get solved and the problem of making knowledge work productive won’t be solved until someone or some group of people assumes the responsibility for doing so.

Just as “Patriotism is the last refuge of a scoundrel,” analysis is the last refuge of the irresponsible. We can’t afford to avoid the accountability for making knowledge work productive by engaging in endless analysis of the problem.

It is a matter of ordinary observation that the smallest number of parts into which a whole can be broken for purposes of analysis is two. I refer to this as “the basic binary split.” The simplest way of making a basic binary split is to affix a label to some portion of a whole, to set it apart from the rest. “They,” for example, implies “not they.” As we talk about knowledge workers throughout the remainder of this article, it is important to keep in mind that *they* are *us*. *We* are the knowledge workers, not some remote, unknown group of other people. As the possum philosopher Pogo once said: “We has met the enemy and they is us.”

Having warned of the dangers of “the basic binary split,” let me immediately make one. Some of the differences between knowledge work and manual work are contrasted in the table below. It is worth pointing out that they reflect not so much the differences between the two kinds of workers as they do the differences between their work.

Characteristics	Manual Work	Knowledge Work
<i>Work Content</i>	Materials	Information
<i>Visibility of Working</i>	High	Low
<i>Causal Linkages</i>	Usually direct	Usually indirect
<i>Time Frames</i>	Usually immediate	Often delayed
<i>Results</i>	Tangible	Tangible & Intangible
<i>Work Routines</i>	Prefigured	Configured
<i>Thinking Mode</i>	Linear	Non-Linear
<i>Locus of Control</i>	Management	The Worker
<i>Role of the Worker</i>	Instrument	Agent

The Current Focus

Making knowledge work productive is itself a form of knowledge work. If the work of making knowledge work productive is to be productive, it must focus on the right issues as opposed to being unfocused or focused on the wrong issues. Currently, I believe there is in some quarters a lack of focus and in others the wrong focus. Both are unproductive.

The lack of focus has to do with what appears to be a myth about the nature of knowledge work; namely, that it is “invisible.” If it is indeed invisible, then attempts to control it are doomed to failure because it cannot be observed. This mistaken belief is sustained by a failure to distinguish between work and working. Dale E. Zand, for example, had this to say about knowledge work in his book, *Information, Organization, and Power*: “The essence of the knowledge organization is that work is done in the head. This means that knowledge work can’t be seen.” Zand’s first point is accurate: knowledge work is done in the head. But doing it is called working, not work. Thus, his second point is incorrect; it is working that has become invisible, not work. As Peter Drucker states in *Management*, “Even the most intangible piece of work is outside the worker and independent of him [or her].”

Work is what the worker does and working is the activity of doing it. The work of a medical claims examiner, for instance, consists of a set of information-processing operations that, for the most part, can be expressed quite visibly in the form of algorithms and decision tables. The work consists of these operations; working consists of carrying them out; and the worker, of course, is the claims examiner. The task of making knowledge work productive is not helped at all by the belief that it is “invisible.”

Not so long ago, most of the work in this country was manual work, that is, it was materials-based. The activity of working was primarily physical and thus observable, as it had been for thousands of years. Control over work was equated with control over the process of work and that in turn was equated with control over the activity of the worker. Conditions were such that, as John Kenneth Galbraith noted of working life on the plantation toward the end of those millennia, “The laggard worker could easily be identified. And he [or she] could then be encouraged to greater productivity by the voice of the overseer and his whip.” The focus of control was clearly on the worker and the locus of control rested with the overseer.

“Overseers” have since given way to supervisors but the basic function remains the same: To “see over” the work. With the sudden disappearance from view of the activity of working, this might prove to be a formidable task. As Shoshana Zuboff, a social psychologist at the Harvard Business School, is quoted as saying in *Psychology Today*, “It is much easier to envision how to exert managerial control over a set of people turning bolts and screws than it is to envision such control over people who must mentally attend to and process information.”

Ms. Zuboff’s point is well taken but, like Zand’s, it is off the mark. The proper focus of a work control system is the work itself, not the worker. Control over the worker was merely a convenient means to the end of being able to control the work, a means now denied us by the shift to knowledge work. Perhaps it is the loss of this convenient control over work that obscures the simple conclusion that if working is indeed under the control of the worker, then it is the worker who must make his or her work productive.

New means of controlling the work of our organizations and institutions are needed – and fast. Moreover, knowledge workers will have to develop them. No one else can do it and it’s got to be done. It’s also got to be done well. As Drucker notes, “Knowledge work, therefore, needs far better design, precisely because it cannot be designed for the worker. It can be designed only by the worker.”

Revamping our work control systems carries with it the risk of not getting it right, of control systems gone wrong. The risk of control systems gone wrong is not to the worker but to the employing organization. This is especially true in the case of that control system known as “policy.”

In their zeal to import and exploit new technologies, organizations will sometimes prematurely “institutionalize” a new technology in the form of policy. Technology as policy as control goes wrong for the following reasons. First, the new policy is inevitably written in the “jargon” of the technologists. Second, there are never enough technologists to go around so, instead of implementing the technology they wind up on staff assignments where their job is to interpret the new policy for all the non-technologists who have to do the work that is subject to the policy. Installing and interpreting policy is almost as good as making it. In other words, policy written in the jargon of technology creates a tremendous transfer of power to the technologists. Worse, it leads to no small amount of wasted energy.

Energy is wasted in at least three ways as a result of policy gone wrong. First, non-technologists expend considerable energy trying to comply with a policy they don’t understand. Second, those who find themselves at odds with the policy expend a great deal of energy fighting it and thus incur corresponding energy costs on the part of its defenders. Third, the technologists in the organization are diverted from their true task, which is the application of the technology to the organization’s benefit and instead become “translators” and then “checkers” who check on compliance. They are diverted from performing to policing.

Two examples of this are AT&T’s institutionalization of instructional technology in the form of its *Training Development Standards* (TDS) and the military’s institutionalization of the same technology in its model of instructional systems development (ISD).

By their very nature, process controls such as TDS and ISD, specify a process, one that from its technologist authors’ point of view is often seen as *the* process. Naturally, the prescribed process excludes other approaches. If it is not exclusionary in nature then it is not policy and can be safely ignored. Policy acts to reduce and restrain options. Therein lies another risk to the organization: inappropriate restraint.

Knowledge workers, by virtue of the mental and verbal nature of their activities, work in social situations. Of these, Drucker said, “It is almost impossible to prefigure the responses appropriate to a given situation.” Yet, through policy, we continue our attempts to impose “the one right way,” a relic of the world of manual work.

The current focus, then, is misplaced. It belongs on work, not the worker, and on expanding options for doing it, not restricting them. To refocus on those things that will enable us to make knowledge work more productive will require us to change that which governs our focus: our mind set, our world view, our frame of reference.

As stated earlier, working is under control of the worker. How then do workers control working? Surprisingly, the answer is they don't. Instead, the workers' actions control their perceptions of the world about them, including the work they do. All people, including workers, act in ways that keep their perceptions of matters aligned with the way they want things to be. The way people want things to be is known by many names: goals, objectives, targets, aims, intentions, etc. Or, as described in William T. Powers' book, *Behavior: The Control of Perception*,³ they are known as "reference conditions." If our perceptions indicate a difference between our perceptions of the conditions we seek to control and our reference level for those conditions, we act (i.e., behave) so as to make our perceptions of conditions consistent with our reference levels for those conditions. Clearly, our behavior is a means to the ends we seek. We often seek to obtain the same or similar ends under very different conditions. Accordingly, we must vary our behavior to do so. This entails the exercise of initiative, imagination and no small amount of problem solving. These are not the kinds of demands that have historically been made of manual workers. If we are to change the way we approach work, workers, working and productivity, we must change our view of the worker – and because we are all workers, it is also our view of ourselves.

The Worker as Agent

Our view of the worker has been that of an instrument, a tool to be employed by others. This view of the worker is in turn rooted in our view of organization. Organization is our primary means of realizing the ends of our collective endeavors. Any endeavor – personal, professional, profit-oriented or political – exists in two states: In the abstract, as conception; and in the concrete, as action and its effects. Linking the two is organization.

For several thousand years, the action side of collective human endeavor has relied on physical, observable working activities. As Moïshe Davidowitz, a futurist observed, "For about the first 90,000 years we used each others' muscles. Then we got smart. For about the next 10,000 years we used the muscles of other animals. Then came the industrial revolution and we began using the 'muscles' of machines. For the last 30 years³, we have been using the 'minds' of machines [the computer]."

It was perhaps only natural that a mind-muscle analogue would come to dominate our thinking about organization. This analogue led us to look at organization as an arrangement among people wherein one group serves as the "mind" and another group serves as its "muscles." Connecting the two was the organization's "central nervous system": its chain-of-command. Given this view of organization, control was seen as a matter of ensuring compliance of the "muscles" with the dictates of the "mind." The worker was an instrument of central authority: a tool.

That the worker has been viewed as an instrument or a machine needs no more confirmation than the simple observation that we have replaced many of them with machines, even some of the knowledge workers.

When work was materials-based, both it and working could be subjected to an engineering approach, that is, the most economically efficient work process and working activities could be determined, described and prescribed – and deviations from them proscribed. The worker could be instructed and instructions could be enforced.

Knowledge work, however, is information-based, not materials-based. And, although work itself is still visible, or can be made so, working is no longer visible. Moreover, as was noted earlier, the activities of most knowledge workers must be configured to meet situational demands instead of carried out as unvarying, prefigured routines. In order to be productive, knowledge workers must exercise control over their working activities. This means the control problem facing organizations and institutions is one of obtaining contributions instead of ensuring compliance. The role of knowledge workers is one of agents, not instruments; they are people who act on behalf of their organization, not simply at its behest.

Drucker noted of people in general that, "The human being works best at a configuration of operations rather than at a single operation." He noted of knowledge workers in particular that they "...work best if capable of varying both speed and rhythm fairly frequently." To do this, the knowledge worker must have control over his or her time. Too many knowledge workers are unnecessarily shackled to the work place and its clock. An "honest day's work" for a knowledge worker does not necessarily result from spending eight hours on the premises, especially if any significant portion of that time is spent in "busy work."

It is doubtful if people can sustain a concentrated mental effort in the same way they can a physical one. Boredom sets in and energy levels sag: attention wanders and productivity goes down owing to a lack of concentrated, channeled energy expenditures. It might be the case that the seemingly fragmented nature of managerial work is actually a way of countering low productivity. By working on several different projects during the same time frame, managers can sustain their interest, attention and energy levels. It also seems that sustained mental effort benefits from a sense of excitement. This is lucidly illustrated in John Tracy Kidder's account of the development of a new computer at Data General.⁴ See his book, *Soul of a New Machine*.

³ Davidowitz's remarks were made in 1978.

⁴ Data General has long since disappeared from the scene but Kidder's book is still first-rate reading.

But our view of the worker must change for another reason as well. There is a clear and present danger in a continued view of the worker as instrument. As such, he or she is the obvious focus of control. As B. F. Skinner pointed out in *About Behaviorism*, "Those who are so controlled then take action." He adds, "In other words, they oppose control with counter-control." Perhaps the best-known means of counter-control exercised by workers has been the union.

Many knowledge workers are already organized in unions. Teachers, for instance, have been organized for some time. So have many technical and clerical workers. For the most part, however, managerial and professional workers have failed to organize except in loose, splintered associations. Yet, some professional associations can be extremely powerful, as is the case with the American Bar Association and the American Medical Association. There is some concern that loyalty to these professional associations might exceed loyalty to the employing organization or institution. This is not an ill-founded concern. If knowledge workers are left with no way of exercising control over their destiny except to organize, then organize they will.

Imagine, if you will, a new union: *The United Mind Workers of America*. The very sound of it conjures up images of John L. Lewis and that's enough to give many older managers heartburn. Such a union could represent a revolutionary redistribution of power, one that could shake the very foundations of our society, not to mention its organizations and institutions. Moreover, it could lead to internecine warfare between one class of knowledge worker and another (e.g., between managers and professionals). Worse yet, managers might join. After all, they, too, are human beings who seek control of their destiny.

Fortunately, or unfortunately, depending on your point of view, knowledge workers are unlikely to organize any time soon. We tend to set ourselves apart from one another based on our occupation instead of banding together based on the nature of our work.

Our view of work, especially of knowledge work, will also have to change.

Three Dimensions of Knowledge Work

I have found it useful to distinguish among three dimensions of knowledge work just as I have found it useful to distinguish among three dimensions of working. (More on the differences in working a little later.) The three dimensions of knowledge work, shown in Figure 1, are information processing (IP), knowledge application (KA) and knowledge creation (KC).

Three Dimensions of Knowledge Work




Figure 1

Information processing, as the term implies, is concerned with information processing operations, often routine and repetitive ones. Filling out forms is an example of the information-processing dimension of work. These operations are often routine and repetitive. A good example is the work of a claims examiner. In systems terms, the input is information and the output is information. The work process, therefore, is essentially one of converting information from one form to another (e.g., from a claim submitted to a claim paid or denied). The processor, however, is a human being, not a computer. Knowledge application is closely akin to Jurgen Ruesch's concept of "knowledge in action," which he defines as "...general knowledge that has been flagged as to time, place, person and situation."

Knowledge itself has been defined by Daniel Bell, author of *The Coming of Post-Industrial Society*, as “An organized set of statements of facts or ideas, presenting a reasoned judgment or experimental results, that is transmitted to others through some communication medium in some systematic form.” Peter Drucker defined it more tersely as “... the systematic organization of information and concepts.” Perhaps the most recognizable form of knowledge application work is known to managers as “problem solving,” to physicians as “diagnosis” and to technicians as “troubleshooting.” All three labels refer to the same basic work: applying knowledge to specific situations for the specific purpose of figuring out what to do. Knowledge creation is a form of intelligence work. It is not unlike Edwin Yoder, Jr.’s definition of journalism, that is, “It is a business of finding themes and patterns, many of them arbitrary, that seem to make sense of events.” A case in point is John Naisbitt’s book, *Megatrends*. The sense that he makes of events, and which he presents in his book, is based on a method known as “content analysis,” a simplified version of which is an essential tool in any knowledge worker’s toolbox. Intelligence work or knowledge creation is at the heart of several organizational functions (e.g., marketing, strategic planning, competitive analysis and even problem solving) and it shows up in several managerial tasks (e.g., interviewing, appraising, estimating, planning and politicking).

Just as work should not be confused with working, neither should it be confused with “job” or “occupation.” The work of knowledge workers involves all three dimensions of knowledge work, albeit to varying degrees. Figure 1 illustrates this notion; it shows the three dimensions of knowledge work for a claims examiner, a manager and a researcher. The claims examiner is high on information processing with low requirements for knowledge application and knowledge creation. The manager faces significant requirements on all three dimensions. The researcher has little in the way of information processing requirements and focuses instead on knowledge application and creation.

None of the representations in Figure 1 are meant to describe people. They are instead a way of looking at the multi-dimensional demands made of people by the nature of their work and working activities. The nature of these demands and the form that energy expenditure takes in satisfying them provide significant clues for use in solving the puzzle of making knowledge work more productive. Salespeople, for example, expend a great deal of energy in verbal activity; one of their primary tools is language. Managers expend a great deal of time in mental activity, for instance, in planning, problem solving and decision-making. Some of their primary tools are the mental models they use to decide what is and isn’t relevant to a given situation.

Models and language, these are the tools of the knowledge worker. We cannot be productive without good tools. But we cannot select or use tools without first understanding the nature of our work and the activity of working. We’ve looked briefly at three dimensions of work. Let us now turn our attention to three dimensions of working.

Three Dimensions of Working

Working, as stated earlier, refers to the activity of the worker. This activity can be primarily physical, verbal or mental; however, most jobs entail activities of all three kinds. Figure 2 below presents some representative energy expenditures along these three dimensions. These are based on my recollections as someone who held the jobs in question, not on data collected from others or from any formal research.

Pin-boy was my first paid job. It is a now defunct job category; pin-boys have long since been replaced by automated pin-setting equipment. As Figure 2 shows, the working activities of a pin-boy were primarily physical. The job required very little verbal activity and not much more mental activity (mostly a matter of remembering which frame the bowler was in and deciding which slots in the pin-setting rack could have pins placed in them and which could not).

Sales representative – telephone sales rep to be more precise – was a very different kind of work. Almost all of my energy expenditures as a sales rep were verbal or mental. Selling over the telephone is a two-level kind of activity: on one level, you’re trying to track and influence the buyer and, on the other, you’re trying to develop ways of overcoming objections, etc. The physical component, although slight, was critical. It consisted of filling out order forms. Errors on these forms (e.g., the wrong street address), were costly in more ways than one.

As the head credit clerk behind the counter in the “flagship” store of a retail appliance chain, my work consisted mostly of verbal and mental activity (e.g., dealing with customers who wanted to apply for financing of a major purchase, collecting sales reports from other stores, operating a fax machine and preparing sales summaries by store and by salesperson). The physical component – filling out forms, preparing reports and carrying purchases to the customer’s car – was greater than had been the case as a telephone sales representative.

My work as a technician was the most all-around demanding. As a Fire Control Technician (FT) in the Navy, my work included operating, maintaining and repairing complex, shipboard weapons control systems. It also included a stint in a weaponry research laboratory, working with engineers from organizations such as Vitro, Sperry and the Johns Hopkins Physical Sciences Laboratory. Later, as a Chief Petty Officer and acting division officer, my work also included no small amount of working with people and with paper. All three dimensions – physical, verbal and mental – were called into play in a significant way.

Assigned as an instructor at the Navy’s Instructor Training School in San Diego, I eventually found myself in charge of the Programmed Instruction (PI) Writer’s Course. Developing programmed instructional materials is primarily a verbal and mental activity – writing and analyses. It also included a great deal of interaction with trainees, reference materials as well as a moderate amount of administrative work. As Figure 2 shows, this work was demanding verbally and mentally but it was low with respect to physical demands.

Three Dimensions of Working

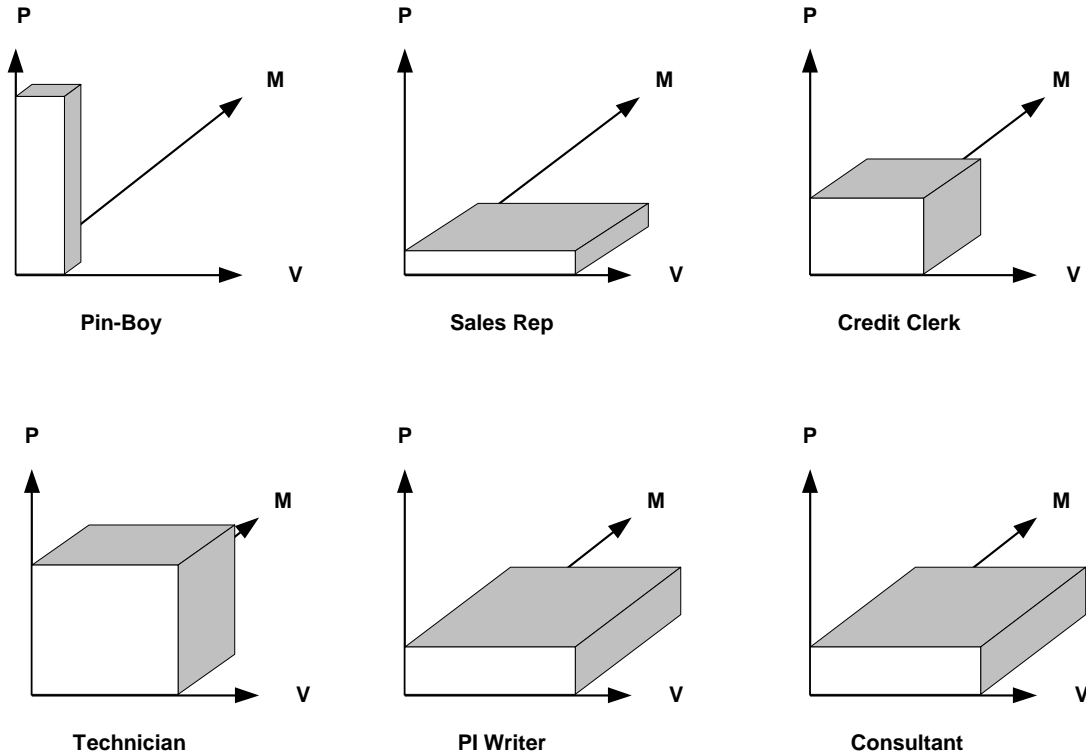


Figure 2

As a consultant and project manager, I have found the major demands to be mental and verbal, more so than in any other kind of work, and with a surprising, although still modest, amount of physical activity. Hours spent at the keyboard require physical effort as well as mental and verbal.

The point of these little “work boxes” is not to inform you of my work history but to illustrate how the activity of working varies along the three dimensions cited. The so-called “manual worker” has a mental and verbal dimension to his or her work just as the so-called “knowledge worker” has a physical dimension to his or hers. Thinking about working activities along these three dimensions directs one’s attention to the nature of the energy expenditures made by the worker in the course of working. To increase the productivity of work, the energy expenditures along all three dimensions must be made more efficient and effective.

With the shift to knowledge work, interest in working tends to focus on two of the three dimensions: mental and verbal. Working, whether mental, physical or verbal, consists of expending energy in a purposeful way over time. Work, of course, consists of the effects or results brought about by these energy expenditures.

To make knowledge work productive, we must concentrate and channel energy along productive lines. Enabling us to do so with respect to manual work was the great contribution of Frederick Winslow Taylor. As yet, we do not have this capability with respect to knowledge work.

Productivity and Knowledge Work

Peter Drucker, chief chronicler of the shift to knowledge work, wrote, "To make knowledge work productive will be the great management task of this century just as to make manual work productive was the great management task of the last century." Some implicit steps in making knowledge work productive have already been discussed:

- Distinguish among work, worker and working
- Focus control on the work, not the worker
- Acknowledge the worker's control over working

Productivity is almost as elusive as knowledge work. Perhaps this is because productivity, like knowledge work, is a constructed reality instead of an observed one. In any case, definitions of productivity are as plentiful as consultants in an economic upturn. Yet, there is no single, agreed upon definition of productivity. This might be because the term is often applied to the nation, segments of its economy, organizations within these segments and to functions and people within organizations. It is small wonder that our definitions sometimes fail us.

Given the information-base of knowledge work, "output" becomes an elusive measure of work. In a monograph titled *Knowledge Worker Productivity*, Ira Gregerman put it this way: "Since it is difficult to quantify the output of knowledge workers, it is difficult to measure their productivity with the classical methods." With knowledge workers, one cannot count the number of widgets coming off the assembly line.

Working, whether mental, verbal or physical, involves expending energy in purposeful ways over time. The work accomplished is more or less productive depending on how well the energy expenditures are concentrated and channeled.

Knowledge workers can be viewed as energy-based work systems having two forms of energy input: (1) the energy provided by the worker and (2) the energy provided by the information, knowledge or intelligence to be processed. It is important to note that the knowledge worker is both generator and processor. Making knowledge work productive entails finding ways and means by which a knowledge worker can become a more effective and efficient producer of knowledge and information as well as a user or consumer.

With an energy-based view of work, improving its productivity becomes clear enough for action purposes, if not for measurement. Where is energy being wasted in work or working? How can energy levels be raised? How can they be sustained? These kinds of questions direct our attention and focus our attention on making knowledge work productive.

With all that has thus far been presented in mind as background, let's examine three instances in which knowledge work was made more productive. The first example illustrates information processing work, the second illustrates knowledge application and the third illustrates intelligence work.

Making Information Work Productive

Many organizations process information from forms via computer (e.g., medical claims, claims for lost or stolen travelers checks, loan and insurance applications, and applications for financial aid). When the computer program is unable to process the information provided, the claim or application is suspended for review by a human being, often known as an "examiner." In three different organizations, a major health insurer, a large financial institution and a leading educational services organization, the work of these examiners was made more productive.

First, the work itself was made visible in the form of algorithms. As used here, an "algorithm" refers to a simplified decision flowchart that depicts the information processing operations the examiners were expected to carry out.

The starting point for developing the algorithms was the physical component of the examiners' work (i.e., the entries made on the suspended documents in order to re-enter them into automated processing). These defined the range of resolution actions they could make.

Given the range of resolution entries for a given suspension, we then identified all the factors that might lead an examiner to settle on a given resolution. These, in turn, were analyzed to determine the kinds and sources of information needed to make them.

Finally, these information sources and requirements, the decisions they enabled, and the entries on the form signifying they had been made, were arrayed in simplified flowchart form: an "algorithm."

As is the case with the child who takes apart a watch, taking apart the information processing operations of the work done by the examiners was a much easier task than the one of reassembling them in the form of algorithms that would handle all possible suspensions. As a matter of fact, in one case, we couldn't. But, the one instance that could not be reduced to an algorithm accounted for less than half of one percent of the suspense stream.

Second, new tools were developed for the examiners based on the algorithms. This was accomplished by preparing prose to accompany the algorithms, augmenting the prose and algorithms with a few decision tables, and packaging the algorithms, prose and decision tables in the form of a desk-level reference manual – a job aid. This manual or job aid communicates to the examiners the operations they are expected to use their heads to perform. It serves to concentrate and channel their energies along productive lines.

Third, training for the examiners was overhauled. That consisted primarily of shortening it and focusing it on using the new tools to resolve suspended forms.

Fourth, in concert with managers and supervisors, new performance and productivity standards were put in place, along with improved feedback mechanisms made possible by the algorithms (e.g., more precise information about where an examiner was going wrong in the resolution of a particular suspension). In short, expectations were crystallized and raised. They were also linked to appraisal and a link to compensation was recommended.

The long and the short of all three of these efforts is that fewer people are able to do more and better work in less time and at less cost. No matter your definition of productivity, that constitutes improvement. The specific nature of the gains varied from company to company and the gains are too extensive to list here in their entirety. However, a few examples will serve to illustrate.

In one company, the use of the algorithms yielded much higher levels of job performance at the end of training. This made possible a 50% reduction in what was known as “100% checking (i.e., having a supervisor inspect every piece of work done by a worker until he or she reached a given level of proficiency). In this same company, it also proved possible to eliminate look-up time, that is, time spent by the examiner in determining the reason for the suspension. To eliminate this unnecessary step, the computer-based processing system was modified so that when a form was suspended, the reason for its suspension was printed on the document. This alone represented a cost savings of approximately \$60,000 per year.

In another company, after the work had been made visible in the form of algorithms, the percentage of the examiners’ time that actually went into performing this work at their workstations was examined. As it turned out, they spent almost 40% of their time away from their workstations.

Roughly one-third of the time away from station was spent standing in line at one centrally located, high-speed copying machine. This machine had been “sold in” on the basis of savings in copy costs. These, however, were not nearly enough to offset the salary costs of the examiners waiting to use the machine and so smaller machines were reintroduced and distributed among the examiners’ workstations.

Roughly another third of the examiners’ time away from station was spent seeking their supervisor’s approval for the resolution of very minor claims. Given that there was no known instance of a supervisor ever disapproving a resolution proposed by an examiner, increasing the examiners’ authorization level further reduced energy expenditures not focused on their work.

Other benefits included reducing backlogs, more evenly distributed workloads, reducing staffing levels and raising standards. In two of the three projects, improvements were made to the computer-based system as a result of making visible the work of the people-based system. In all three projects, inefficiencies in the work itself were identified and removed, making possible higher levels of productivity. Overall increases ranged from 20% to 50%.

The new tools were well received by the workers. One reason for this is that they were intensely involved in developing them. As the project manager on all three projects, I made it clear to all concerned that only the workers knew how to make their work more productive and that I was there to help them do that, not do it for them or in spite of them. Their reactions to their new tools were positive, then, because they helped create them. A couple of the workers’ comments are most revealing:

- “It [the new manual] makes my job easier.”
- “They [the algorithms] let me know exactly what I should be doing.”

In other words, the new tools reduced the ambiguity and uncertainty associated with their work, enabling them to concentrate and channel their energy more productively. The examiners were as pleased as their management that they would no longer be wasting so much of their time and energy.

Not all knowledge work reduces to a set of algorithms. Knowledge application, for example, is frequently far too ambiguous and uncertain to be treated in such a manner. Nevertheless, it, too, can be made more productive. An example from AT&T will illustrate how.

Improving the Productivity of Problem Solving

AT&T, in the course of assessing the training needs of its first, second and third level managers, a population then well in excess of 200,000 people,⁵ determined that problem solving was an area where training was warranted. Given the size of the audience, AT&T decided to develop a proprietary training course instead of using “off the shelf” or vendor-owned programs. My firm was engaged to develop it. Our approach was a deliberate effort to make the problem solving activities of Bell System managers more productive through training. The course we developed, *Managing the Problem Solving Process*, did just that.

The reasoning behind the design of the course is more relevant than the mechanics of its development so it is this “behind the scenes” work that I will review here.

Our initial thinking was about payoffs. We reasoned that if the course were to have any payoffs, they would be felt in the form of reductions in the costs of the problems and in the costs of solving them. These reduced costs could

⁵ The project in question began in 1979, long before AT&T was faced with divestiture.

be realized in one or both of two ways: (1) by coming up with better solutions and (2) coming up with them more quickly. The aim of the course quickly became “better solutions faster.”

We concluded that coming up with better solutions would depend on introducing better information into the problem solving process and that coming up with solutions more quickly would depend on streamlining the information-processing aspects of the problem solving process. In short, we were ourselves being guided by the fundamental notion that our objective was to find ways of concentrating and channeling human and organizational energy.

Information is introduced into the problem solving process by the people who are part of it. In turn, the information they admit to the process is a matter of what they view as relevant. This relevance decision is affected by several factors; for example, the label placed on the problem – which, in turn, determines the frame of reference from which they examine the problem, thereby generating certain kinds of information and not others – and, quite frankly, the kinds of matters that are appropriate to discuss and those that are not.

Several possibilities for increasing the quantity and the quality of information brought to bear in problem solving were thus identified; for example, deliberately changing the label on the problem, changing the frame of reference of the problem solvers, involving more people and thus more frames of reference, and indicating and tracking the kinds of information that are required to solve a problem.

The efficient processing of the information introduced into the problem solving process is a function of the mechanisms used to process it. As Charles Kepner and Benjamin Tregoe noted in their classic little book, *The Rational Manager*, people tend to jump around from problem to solution to cause and back again when discussing problems and often in no particular order. Kepner and Tregoe saw this as extremely inefficient and went on to propose a linear, rational approach to problem solving, one that is essentially a fault isolation or troubleshooting approach and that has become the “bread and butter” of more than one organization’s management training and development program.

Unlike Kepner and Tregoe, we reasoned that the seemingly random generation of information is an inherent characteristic of human beings seeking to make sense out of a situation, especially when working in groups, and so we decided on a strategy of capturing and organizing the information generated instead of attempting to force people to adopt a step-by-step, linear approach. In keeping with this reasoning, it was decided that the training course would focus on managing the problem solving *process* as opposed to managing its *content*. The project team’s assessment was that the AT&T managers, like most people, were already reasonably effective problem solvers and that the primary contribution of the training course would be one of helping them become more efficient (i.e., get there faster). The course provides the managers who attend it with three tools:

1. A conceptual view of the problem solving process.
2. A nine-cell “Problem Information Matrix” that specifies the kind of information required and also serves to capture and organize the information that is generated.
3. Checklists and other job aids intended for a “when you need them” level of usage.

The primary tool is the “Problem Information Matrix” because it reveals what is required in the way of information, what is possessed, and thus the nature of any shortages. It does what any good knowledge work tool should do: it tells you what you need to know that you don’t know.

Managing the Problem Solving Process treats problem solving as an action-oriented and solution-focused process instead of a problem-focused and analysis-oriented activity. A problem is defined as “a situation requiring action.” A solution is defined as “an action that reduces or eliminates the need for further action.” And problem solving is defined as “the search for a solution.”

From a decision-making perspective, the problem solving process is marked by four basic, binary decisions:

1. Do you have a problem?
2. Is it worth investigating?
3. Do you know what to do?
4. Is it worth doing?

Making the aforementioned decisions requires certain kinds of information from various sources. Therefore, we treated problem solving as a “cover the bases” activity, that is, as a matter of ensuring the availability of adequate information of each kind on the “Problem Information Matrix.” The matrix, then, allows the problem solvers to generate information naturally, in a non-linear fashion, and it actually encourages “jumping around” (i.e., of moving from “base” to “base”) until such time as all the bases have been adequately covered. Its effect is to speed up the generating and processing of problem and solution-related information instead of slowing it down.

A formal evaluation of *Managing the Problem Solving Process* was conducted by an independent consulting firm hired by AT&T. The evaluation revealed that the course was indeed meeting its objectives. Another measure of its value was that it quickly migrated from the second-level management curriculum for which it had been initially prepared to the curricula for the first and third levels as well. It also migrated into Quality of Work Life (QWL) and Quality

Circle (QC) applications as well, where it was noted for its ability to facilitate communications regarding cross-functional and cross-divisional problems.

The basis of the success of the course is neither magical nor mystical. Simply put, it helped people more efficiently and effectively generate, organize and process information, knowledge and intelligence about problems. In so doing, it helped them improve the productivity of one particular kind of knowledge work: problem solving.

As this example illustrates, making problem solving work productive is not nearly as systematic and rigorous a process as was the case with the work of the examiners. In part, this owes to higher levels of ambiguity and uncertainty inherent in problem solving and, in part, it owes to the fact that influencing the workers' standards for their own working activities was a different undertaking in the two cases. Influencing the examiners' standards was a fairly straightforward matter. They could literally "see" what we were proposing and how it would pay off. Influencing the managers who were attending the problem solving training relied a bit more on persuasion, trust and experimentation. There, the proof of the productivity pudding was very much in the eating. But, once they tried it, they liked it.

Matters get even fuzzier when you move beyond information processing and knowledge application kinds of work and move into what I think of as "intelligence work." Nevertheless, it, too, can be made more productive.

Making Intelligence Work Productive

There is a point in all workers' employment, especially managers and other professionals, where a simple form of intelligence work is critical; namely, in the first few days and weeks, or perhaps months, in a new assignment. "Getting up to speed" in a new assignment is a classic form of putting together isolated bits and pieces of information in order to arrive at some well-reasoned, sound conclusions about "the way things work around here." No one hands you this information on a silver platter; you have to find it out for yourself. Moreover, this requirement repeats itself with each new assignment. Many organizations are accustomed to accepting reduced levels of performance and productivity from employees in the early stages of new assignments. This is often referred to as "the honeymoon period." This does not have to be the case. The learning curve on the front end of a new assignment can be systematically and dramatically shortened.

When AT&T was confronted with divestiture and its associated massive restructuring, management was particularly interested in reducing the learning curve on the front end of new assignments. An effort to shorten that learning curve was made and, by all accounts, it succeeded. It was a simple, low-cost effort. All it entailed was putting together an initial "laundry list" of the kinds of things a new job incumbent needs to investigate – regardless of the specific job. The list ranged from things as mundane as locating the copying machine to "psyching out" the new boss' communication style and preferences to analyzing the distribution of power and authority in one's new organization. People taking new assignments were identified and asked to "de-bug" the laundry list. Based on their feedback, it was revised and then published by AT&T as a small book titled *New Assignment – New Challenge*.

Follow-up interviews with newly assigned managers and their bosses indicated that people who used the book did indeed get up to speed more quickly than would otherwise have been the case. One user reported that her boss told her after their initial meeting that "His entire picture of her had changed" as a result of one question she asked from the book: "What, if anything around here, is in danger of falling through the cracks?"

Some senior, more experienced managers viewed the book as too simplistic. However, it wasn't intended for them. Presumably, they were quite accustomed to getting up to speed in new assignments and quite successful at it. For less experienced people, the book was a useful tool. My favorite assessment of the book comes from an experienced middle manager who, upon reviewing the book for us, exclaimed, "My God, you've written it all down!" In one sense, the book might have been a primitive example of what Daniel Bell meant by "the codification of theoretical knowledge."

Making intelligence work productive is itself an instance of intelligence work. I must confess to an inability to approach it in a procedural way. But, as Peter Drucker has said:

"We need to know much more than we now know about the dimensions of working and their relationships. We are dealing with a configuration likely to defy analysis. Nevertheless, the manager has to manage now. He has to find solutions – or at least accommodations – which will enable him to make work productive and the worker achieving. He has to understand what the demands are. He cannot expect to succeed by continuing the practices of the last two hundred years. He will have to develop new approaches, new principles, and new methods – and fast."

Some Tips on Making Knowledge Work Productive

In my experience, making knowledge work productive is like working on any other extremely complex problem: you don't solve it all at once in one fell swoop but, rather, in a dogged, determined "bits and pieces" way. What follows are some ways I have found of improving the productivity of my own work.

Do A Prototype

One way I have found of reducing the ambiguity, uncertainty and risk associated with my work as a consultant is through the use of a prototype effort.

A prototype project is a much scaled-down version of the total project being contemplated. The basic idea is simple enough: “peel off” a piece of the main effort and complete it from start to finish. This allows you to test the feasibility of the approach you intend using, to demonstrate its value, to more accurately gauge the overall project, and to establish important even if preliminary working relationships. If the approach is not going to work, this is discovered long before the bulk of resources are committed and, perhaps, expended.

The length of a prototype effort can range from a few days to a few weeks or more, depending on the size of the main effort behind it. It also has the distinct advantage of making the project itself very productive early on.

In two of the information work projects described earlier, a prototype effort was conducted at the outset. One was five days long; the other was three weeks long. In both cases, we isolated a small piece of the examiners’ work and focused on doing with it what we intended doing with the main body of work, that is, take it apart, develop algorithms and job aids, and demonstrate their viability in improving the volume and accuracy of the examiners’ work. In one case, the size of the main project was revised downward and in the other it was revised upward. In both cases, the feasibility of proceeding was based on first-hand knowledge of what it would take to complete the main project. Risk was reduced.

The hallmarks of the shift to knowledge work are ambiguity and uncertainty. Risk increases with ambiguity and uncertainty and risk acts to constrain action. Yet, there is no work without action. A prototype increases productivity because it decreases the ambiguity and uncertainty surrounding the work.

Avoid Specific Task Assignments

Another way I have found of increasing the productivity of project team members’ work (not to mention my own) is to avoid specific task assignments. This is a response to the impossibility of exercising direct supervision over the working activities of a group of knowledge workers. People assigned to a project team should be viewed, treated and behave as peers. All share equally in the accountability for seeing the project through to its completion.

Somehow, if everyone is keeping track of everything, nothing seems to fall through the cracks. At the last minute, someone usually says something like, “Uh, who’s going to take care of ...?” and someone does.

Although this clearly flies in the face of conventional wisdom, it makes a great deal of sense in terms of managing people who must manage themselves and who must see to making their own work productive. It relocates the pressure to perform between peers instead of between boss and subordinate and it forces initiative on the part of the team members. In short, people keep themselves productively employed on tasks that all need doing as opposed to having a boss assign tasks to them so as to keep them busy.

Isolate & Insulate

Consistent with the basic notion of concentrating and channeling energy along productive lines is an old management tactic: “Lock ‘em in a room.” In all of the projects mentioned earlier, the project team “holed up” in one large room, usually referred to as the “bull pen,” and didn’t return to regularly assigned workspaces until the project or prototype was completed.⁶

There are two reasons for this isolation. First, there is the goal of insulating the project team from day-to-day disturbances. Second, and perhaps more important, a “bull pen” or other communal space fosters the sharing of information, knowledge and intelligence among team members.

The absence of specific task assignments forces initiative and shared information enables it.

Exploit Parkinson’s Law

Parkinson’s Law states that work expands to fill the available time. One way of exploiting this law is to collapse the available time. On occasion, shorten instead of extending the deadlines. When this is done, people seem to display a remarkable ability to sort the wheat from the chaff and get on with whatever is important at the time – without discussing or debating priorities in a formal sense. Within the boundaries posed by the new shortened, deadlines, all manner of adjustments, changes and realignments can take place. Shortening the deadlines also reduces the pressure for progress reporting, which is often more a case of talking about progress than making it.

Another way of exploiting Parkinson’s Law is to refuse to extend deadlines until the very last minute – even when extensions are clearly in the cards – and then extend them by only a modest amount. One reason is to keep the pressure on; another is to control “the available time,” to keep it at a minimum.

Exploiting Parkinson’s Law seems to be a very good fit with one of the properties of excellent companies described by Peters & Waterman in their book, *In Search of Excellence*, namely, “Simultaneous Loose-Tight Controls.”

The Fundamental Task

The preceding examples illustrate the fundamental means of making knowledge work productive; namely, concentrating and channeling energy expenditures along productive lines.

⁶ These were all small teams of 4-6 people. A “bull pen” would be impractical for larger teams.

The ambiguous and uncertain nature of knowledge work precludes a thorough, *a priori*, analytical approach to making it productive. It is instead, in my opinion, a “bits and pieces” task, guided by the basic notion that the objective is to increase the amount of energy brought to bear on the work and to focus that energy so that as little as possible of it is wasted. This means being constantly on the alert regarding where energy is going, where it ought to be going, how it is being wasted and what to do about it. In a very real sense, it is as Peters and Waterman said of Amoco’s success: a matter of “drilling more wells.”

Making knowledge work productive is itself a form of intelligence work. This means we have to piece together an approach to it based on small and isolated successes in several areas. In other words, a strategy and a system for improving the productivity of knowledge work is not going to spring forth full-blown but will instead be stitched together much like a patchwork quilt. As a former technician, I am inclined to say we’re in for a long period of jury-rigging things before we can turn them over to the industrial engineers. That, to me, is the real challenge of the shift to knowledge work.

Concluding Remarks

In this article I have tried to spell out the nature of the challenge of the shift to knowledge work as I see it. I have also tried to make clear that part of the challenge is a fundamental redistribution of power in our organizations and institutions. This owes to the fact that the activity of working has moved out of the realm of observable events and into the currently unobservable realm of mental activity. We are faced with a new set of problems never before encountered on this scale. Knowledge workers have been with us since the beginnings of collective human endeavor but they have never before constituted the majority of the work force.

Before we can come up with new ways of dealing with these new problems, we must face two facts: (1) organizations are collections of solutions to old problems and (2) we must shed the mindsets or frames of reference that gave rise to those old solutions if we are to look at these new problems differently.

It has been my objective in this paper to get you, the reader, to look at work, worker, working, productivity and the productivity of knowledge work in ways that might be different from those to which you are accustomed.

As should be apparent, my thinking about ways to improve the productivity of knowledge work has been greatly influenced by the writings and ideas of Peter Ferdinand Drucker. That is as it should be because he has been its chief chronicler, analyst and commentator. Let me conclude his piece of the action with a few of his thoughts on this subject:

- “The social, political, and economic challenges in and from the work force in industrially advanced countries increasingly arise out of the dynamics of knowledge work and out of the needs of the knowledge workers.”
- “...we did not have – and still do not have – much knowledge about managing the knowledge worker and making him productive.”
- “That we do not yet know how to manage knowledge workers for performance is hardly surprising considering how recent the shift to knowledge work has been.”
- “The most important thing we know is that work and working are fundamentally different phenomena.”
- “...we are beginning to apply knowledge to work itself.”

A Final Comment

My last point is perhaps the most obvious. That is why I have saved it until last.

If it is true that the locus of control over working has shifted from management to the worker, then it seems reasonable to conclude that, as knowledge workers, our minds have become targets for influence. Indeed they are.

Some of these efforts are quite overt and “up front,” as when David McClelland seeks to influence achievement by getting people to think in terms of achieving. Others are subtler; for example, the techniques associated with Neuro-Linguistic Programming (NLP). And some are downright surreptitious; for example, much advertising and most propaganda, including that promulgated by management.

Of late I have taken to analyzing the content and presentation of various professional journals to see if I can spot ways in which the writers are subtly trying to shape my impressions and conclusions. Examples abound. One involved the author slipping a reference to his company between two paragraphs describing impressive successes. But, as careful scrutiny revealed, these successes were not attributable to the author’s company. Another instance was marked by the manipulation of statistical data to support a contentious point. A third entailed the “seeding” of articles with offhand references to some mysterious technique of magical powers. The list goes on and on.

The point of the preceding is to remind us all that, as knowledge workers who now control the activity of working and who now own the means of production, we have likely become the targets of efforts aimed at influencing our beliefs, our standards, our values, our very minds. At least some of the information we are fed is quite likely contaminated. In our thirst for knowledge, we must take care lest we drink too often from polluted waters. Yet, we have little choice except to return time and time again to the well because, as Thomas Sowell points out in *Knowledge & Deci-*

sions, "...we are all in the business of selling and buying knowledge from one another, because we are each so profoundly ignorant of what it takes to complete the whole process of which we are a part."

References

1. Bandler, R. and J. Grinder (1975). *The Structure of Magic I & II*. Science and Behavior Books: Palo Alto.
2. Bell, Daniel J. (1973). *The Coming of Post-Industrial Society*. Basic Books: New York.
3. _____ (1979). "Communications technology – for better or for worse." *Harvard Business Review* (May-June): Harvard University: Boston.
4. Davidowitz, M. (1978). "An Invited Address," 4th *Annual Inter-Service Symposium on Human Resources Management and Development*. U. S. Naval Postgraduate School: Monterey.
5. Drucker, P. F. (1954). *The Practice of Management*. Harper & Row: New York.
6. _____ (1968). *The Age of Discontinuity*. Harper & Row: New York.
7. _____ (1973). *Management*. Harper & Row: New York.
8. _____ (1980). *Managing in Turbulent Times*. Pan Books: London.
9. Galbraith, J. K. (1977). *The Age of Uncertainty*. Houghton Mifflin: Boston.
10. Goleman, D. (1983). "The Electronic Rorschach" in *Psychology Today* (February). New York.
11. Gregerman, I. B. (1981). *Knowledge Worker Productivity*. American Management Association: New York.
12. Kepner C. and B. Tregoe (1965). *The Rational Manager*. McGraw-Hill: New York.
13. Kidder, T. (1981). *The Soul of A New Machine*. Avon Books: New York.
14. McClelland, D. (1969). "Toward a Theory of Motive Acquisition" in *The Planning of Change* (2nd Edition), Bennis, Benne and Chin (Editors). Holt, Rinehart & Winston: New York.
15. Naisbitt, J. (1982). *Megatrends*. Warner Books: New York.
16. Nickols, F. W. (1980). *Managing the Problem Solving Process*. AT&T: Basking Ridge.
17. _____ (1982). *New Assignment-New Challenge*. AT&T: Basking Ridge.
18. Peters, T. J. and R. H. Waterman, Jr. (1982). *In Search of Excellence*. Harper & Row: New York.
19. Powers, W. T. (1973). *Behavior: The Control of Perception*. Aldine Publishing Company: Chicago.
20. Ruesch, J. (1975). *Knowledge in Action*. Jason Aronson: New York.
21. Skinner, B. F. (1974). *About Behaviorism*. Alfred A. Knopf: New York.
22. U.S. Department of Commerce. *Historical Statistics of the United States*. Washington, D.C.
23. U.S. Senate (1982). *Hearings*. Senate Subcommittee on Employment and Productivity: Washington, D.C.
24. Yoder, E. M., Jr. (1983). "Two Storytellers" in *The Washington Post* [Op-ed] (March 24): Washington, D.C.
25. Zand, D. E. (1981). *Information, Organization and Power*. McGraw-Hill: New York.

Afterword

For a while, I thought about attempting a version of this article that reflected my thinking some 20 years later. However, as noted earlier, my thinking remains pretty much the same.

It could also be argued that I should update the article to reflect the thinking of others since 1983. To be perfectly honest, I have stayed on top of the knowledge work and knowledge management literature and there is little in it that would change what I have to say. Drucker has remained remarkably consistent during this time and most other authors have written more about knowledge management than knowledge work and how to make it productive. About all that an update based on the writings of others would accomplish is produce a list of references with more recent entries and I don't view that as particularly productive work on my part.

One aspect of my thinking has changed: I have concluded that, in many ways, knowledge work is a myth. There is just work and, although distinctions can be drawn between "manual work" and "knowledge work," I'm much less convinced than once I was that such distinctions are crucial in making what we call knowledge work more productive. The basic issues of making work productive are the same whether one is dealing with manual or knowledge work. In this vein, there are two papers on my articles web site that might be of interest to readers who have made it this far. Links are embedded in the text below.

The first is titled "[Knowledge Work is A Myth](#)." In it, I explore the reasoning behind the title, mainly in terms of four categories of work that can be derived from the interplay of the content of the work (i.e., materials and information) and the work routine (i.e., prefigured or configured).

The second is titled "[Making Work Productive](#)." It is an essay on work and work control systems. It is the product of a symposium on systems theory and thinking in HR and published in 1989 by the symposium's sponsors, ASTD and the University of Minnesota. I have since revised and updated it.

Finally, if you have made it this far, I thank you for your patience and interest and I hope you find at least a few ideas in this paper that you can put to good use in making your own work and the work of others more productive.

Fred Nickols
January 16, 2004

For More Information

Fred Nickols can be contacted by e-mail at nickols@att.net. Other articles of his can be found by going to his web site at www.nickols.us and clicking on the link to articles or by going to www.skullworks.com.