

California Management Review



Knowledge-Worker Productivity:
The Biggest Challenge

Peter F. Drucker

© 1999 by The Regents of
the University of California

University of California
Berkeley
Haas School of Business

Knowledge-Worker Productivity: THE BIGGEST CHALLENGE

Peter F. Drucker

The most important, and indeed the truly unique, contribution of management in the 20th century was the fifty-fold increase in the productivity of the *manual worker* in manufacturing. The most important contribution management needs to make in the 21st century is similarly to increase the productivity of *knowledge work* and *knowledge workers*. The most valuable assets of a 20th-century company was its *production equipment*. The most valuable asset of a 21st-century institution (whether business or non-business) will be its *knowledge workers* and their *productivity*.

The Productivity of the Manual Worker

First, we must take a look at where we are. It was only a little over a hundred years ago that for the first time an educated person actually *looked* at manual work and manual workers, and then began to study both. The Greek poet Hesiod (eighth century B.C.) and the Roman poet Virgil (700 years later) sang about the work of the farmer. Theirs are still among the finest poems in any language, but neither the work they sang about nor their farmers bear even the most remote resemblance to reality, nor were they meant to have any. Neither Hesiod nor Virgil ever held a sickle in their hands, ever herded sheep, or even looked at the people who did either. When Karl Marx, 1900 years after Virgil, came to write about manual work and manual workers, he too never looked at either, nor had he ever as much as touched a machine. The first man to do both—that is, to work as a manual worker and then to study manual work—was Frederick Winslow Taylor (1856-1915).

Throughout history there have been steady advances in what we today call “productivity” (the term itself is barely fifty years old). They were the result

of new tools, new methods, and new technologies; they were advances in what the economist calls “capital.” There were few advances throughout the ages in what the economist calls “labor”—that is, in the productivity of the worker. It was axiomatic throughout history that workers could produce more only by working harder or by working longer hours. The 19th-century economists disagreed about most things as much as economists do today. However, they all agreed—from David Ricardo through Karl Marx—that there are enormous differences in *skill* between workers, but there are none in respect to productivity other than between hard workers and lazy ones, or between physically strong workers and weak ones. Productivity did not exist. It still is an “extraneous factor” and not part of the equation in most contemporary economic theory (e.g., in Keynes, but also in that of the Austrian School).

In the decade after Taylor first looked at work and studied it, the productivity of the manual worker began its unprecedented rise. Since then, it has been going up steadily at the rate of 3% per annum compound—which means it has been risen fifty-fold since Taylor. On this achievement rest *all* of the economic and social gains of the 20th century. The productivity of the manual worker has created what we now call “developed” economies. Before Taylor, there was no such thing—all economies were equally “underdeveloped.” An underdeveloped economy today—or even an “emerging” one—is one that has not, or at least has not yet, made the manual worker more productive.

The Principles of Manual-Work Productivity

Taylor’s principles sound deceptively simple. The first step in making the manual worker more productive is to look at the task and to analyze its constituent motions. The next step is to record each motion, the physical effort it takes, and the time it takes. Then motions that are not needed can be eliminated; and whenever we have looked at manual work, we have found that a great many of the traditionally most-hallowed procedures turn out to be waste and do not add anything. Then, each of the motions that remain as essential to obtaining the finished product is set up so as to be done the simplest way, the easiest way, the way that puts the least physical and mental strain on the operator, and the way that requires the least time. Next, these motions are put together again into a “job” that is in a logical sequence. Finally, the tools needed to do the motions are redesigned. Whenever we have looked at any job—no matter for how many thousands of years it has been performed—we have found that the traditional tools are wrong for the task. This was the case, for instance, with the shovel used to carry sand in a foundry (the first task Taylor studied). It was the wrong shape, the wrong size, and had the wrong handle. We found this to be equally true of the surgeon’s traditional tools. Taylor’s principles sound obvious—effective methods always do. However, it took Taylor twenty years of experimentation to work them out.

Over these last hundred years, there have been countless further changes, revisions, and refinements. The name by which the methodology goes has also changed over the past century. Taylor himself first called his method “Task Analysis” or “Task Management.” Twenty years later it was re-christened “Scientific Management.” Another twenty years later, after the First World War, it came to be known as “Industrial Engineering” in the U.S. and Japan, and as “Rationalization” in Germany.

To proclaim that one’s method “rejects” Taylor or “replaces” him is almost standard “public relations.” For what made Taylor and his method so powerful has also made it unpopular. What Taylor *saw* when he actually looked at work violated everything poets and philosophers had said about work from Hesiod and Virgil to Karl Marx. They all celebrated “skill.” Taylor showed that in manual work there is no such thing. There are only simple, repetitive motions. What makes them more productive is *knowledge*, that is, the way the simple, unskilled motions are put together, organized, and executed. In fact, Taylor was the first person to apply knowledge to work.¹

This also earned Taylor the undying enmity of the labor unions of his time, all of which were craft unions and based on the *mystique* of craft skill and their monopoly on it. Moreover, Taylor advocated—and this is still anathema to a labor union—that workers be paid according to their productivity—that is, for their output, rather than for their input (e.g., for hours worked). However, Taylor’s definition of work as a series of operations also largely explains his rejection by the people who themselves do not do any manual work: the descendants of the poets and philosophers of old, the Literati and Intellectuals. Taylor destroyed the romance of work. Instead of a “noble skill,” it becomes a series of simple motions.

Nevertheless, every method during these past hundred years that has had the slightest success in raising the productivity of manual workers—and with it their real wages—has been based on Taylor’s principles, no matter how loudly his antagonists proclaimed their differences with Taylor. This is true of “work enlargement,” “work enrichment,” and “job rotation”—all of which use Taylor’s methods to lessen the worker’s fatigue and thereby increase the worker’s productivity. It is also true of such extensions of Taylor’s principles of task analysis and industrial engineering as Henry Ford’s assembly line (developed after 1914, when Taylor himself was already sick, old, and retired). It is just as true of the Japanese “Quality Circle,” “Continuous Improvement” (*Kaizen*), and “Just-In-Time Delivery.”

The best example, however, is W. Edward Deming’s “Total Quality Management.” What Deming did—and what makes Total Quality Management effective—is to analyze and organize the job exactly the way Taylor did. However, he also added Quality Control (around 1940) that was based on a statistical theory that was only developed ten years after Taylor’s death. Finally, in the 1970s, Deming substituted closed-circuit television and computer simulation for

Taylor's stopwatch and motion photos. Deming's Quality Control Analysts are the spitting image of Taylor's Efficiency Engineers and function the same way.

Whatever his limitations and shortcomings—and he had many—no other American, not even Henry Ford, has had anything like Taylor's impact. "Scientific Management" (and its successor "Industrial Engineering") is the one American philosophy that has swept the world—more so even than the Constitution and the Federalist Papers. In the past century, there has been only one worldwide philosophy that could compete with Taylor's: namely, Marxism. In the end, Taylor has triumphed over Marx.

During the First World War, Scientific Management swept through the U.S. together with Ford's Taylor-based assembly line. In the 1920s, Scientific Management swept through Western Europe and began to be adopted in Japan.

During the Second World War, both the German achievement and the American achievement were squarely based on applying Taylor's principles to Training. The German General Staff, after having lost the First World War, applied "Rationalization" (i.e., Taylor's Scientific Management) to the job of the soldier and to military training. This enabled Hitler to create a superb fighting machine in the six short years between his coming to power and 1939. In the U.S., the same principles were applied to the training of an industrial work force, first tentatively during the First World War and then, with full power, during the Second World War. This enabled the Americans to outproduce the Germans, even though a larger proportion of the U.S. than the German male population was in uniform and thus not in industrial production. Then, training-based Scientific Management gave the U.S. civilian work force more than twice—if not three times—the productivity of the workers in Hitler's Germany and in Hitler-dominated Europe. Scientific Management thus gave the U.S. the capacity to outnumber both Germans and Japanese on the battlefield and yet still outproduce both by several orders of magnitude.

Since 1950, economic development outside the Western World has largely been based on copying what the U.S. did in the Second World War, i.e., on applying Scientific Management to making the manual worker more productive. All earlier economic development had been based on technological innovation—first in France in the 18th century, then in Great Britain from 1760 until 1850, and finally in the new economic Great Powers, Germany and the U.S., in the second half of the 19th century. The non-Western countries that developed after the Second World War, beginning with Japan, eschewed technological innovation. Instead, they imported the training that the U.S. had developed during the Second World War based on Taylor's principles and they used it to make highly productive, almost overnight, a still largely unskilled and pre-industrial work force. (In Japan, for instance, almost two-thirds of the working population were still, in 1950, living on the land and unskilled in any work except cultivating rice). However, while highly productive, this new work force was still—for a decade or more—paid pre-industrial wages so that these countries—first Japan, then Korea, then Taiwan and Singapore—could produce the

same manufactured products as the developed countries, but at a fraction of their labor costs.

The Future of Manual-Worker Productivity

Taylor's approach was designed for manual work in *manufacturing*, and at first applied only to it. Nevertheless, even within these traditional limitations, Taylor's approach still has enormous scope. It is still going to be the organizing principle in countries in which manual work, and especially manual work in manufacturing, is the growth sector of the society and economy—that is, "Third World" countries with very large and still growing numbers of young people with little education and little skill.

However, there is equal—or even greater—opportunity in the *developed* countries to organize non-manufacturing production (i.e., production work in services) on the production principles now being developed in manufacturing—and that means applying Industrial Engineering to the job and work of the individual service worker. There is equally a tremendous amount of knowledge work—including work requiring highly advanced and thoroughly theoretical knowledge—that includes *manual* operations. The productivity of these operations also requires Industrial Engineering.

Still, in developed countries, the central challenge is no longer to make manual work more productive—after all, we know how to do it. The central challenge will be to make knowledge workers more productive. Knowledge workers are rapidly becoming the largest single group in the work force of every developed country. They may already compose two-fifths of the U.S. work force—and a still smaller but rapidly growing proportion of the work force of all other developed countries. It is on their productivity, above all, that the future prosperity—and indeed the future survival—of the developed economies will increasingly depend.

What We Know About Knowledge-Worker Productivity

Work on the productivity of the knowledge worker has barely begun. In terms of actual work on knowledge-worker productivity, we will be in the year 2000 roughly where we were in the year 1900 in terms of the productivity of the manual worker. Nevertheless, we already know infinitely more about the productivity of the knowledge worker than we did then about that of the manual worker. We even know a good many of the answers. We also know the challenges to which we do not yet know the answers, and on which we need to go to work.

Six major factors determine knowledge-worker productivity.

- Knowledge-worker productivity demands that we ask the question: "*What is the task?*"

- It demands that we impose the responsibility for their productivity on the individual knowledge workers themselves. Knowledge Workers *have* to manage themselves. They have to have *autonomy*.
- Continuing innovation has to be part of the work, the task and the responsibility of knowledge workers.
- Knowledge work requires continuous learning on the part of the knowledge worker, but equally continuous teaching on the part of the knowledge worker.
- Productivity of the knowledge worker is not—at least not primarily—a matter of the *quantity* of output. *Quality* is at least as important.
- Finally, knowledge-worker productivity requires that the knowledge worker is both seen and treated as an “asset” rather than a “cost.” It requires that knowledge workers *want* to work for the organization in preference to all other opportunities.

Each of these requirements (except perhaps the last one) is almost the exact opposite of what is needed to increase the productivity of the manual worker. In manual work, of course, quality also matters. However, lack of quality is a restraint. There has to be a certain minimum quality standard. The achievement of Total Quality Management—that is, of the application of 20th century Statistical Theory to manual work—is the ability to cut (though not entirely to eliminate) production that falls below this minimum standard.

In most knowledge work, quality is not a minimum and a restraint. Quality is the essence of the output. In judging the performance of a teacher, we do not ask how many students there can be in his or her class. We ask how many students learn anything—and that’s a quality question. In appraising the performance of a medical laboratory, the question of how many tests it can run through its machines is quite secondary to the question of how many tests results are valid and reliable. This is true even for the work of the file clerk.

Productivity of knowledge work therefore has to aim first at obtaining quality—and not minimum quality but optimum if not maximum quality. Only then can one ask: “What is the volume, the quantity of work?” This not only means that we approach the task of making more productive the knowledge worker from the quality of the work rather than the quantity, it also means that we will have to learn to define quality.

What Is the Task?

The crucial question in knowledge-worker productivity is: *What is the task?* It is also the one most at odds with manual-worker productivity. In manual work, the key question is always: *How should the work be done?* In manual work, the task is always given. None of the people who work on manual-worker productivity ever asked: “What is the manual worker supposed to do?” Their only question was: “How does the manual worker best do the job?” This was just as true of Frederick W. Taylor’s Scientific Management as it was true of the people

at Sears Roebuck or the Ford Motor Company who first designed the assembly line, and as it is true of W. Edward Deming's Total Quality Control.

Again, in knowledge work the key question is: What is the task? One reason for this is that knowledge work, unlike manual work, does not program the worker. The worker on the automobile assembly line who puts on a wheel is programmed by the simultaneous arrival of the car's chassis on one line and the wheel on the other line. The farmer who plows a field in preparation for planting does not climb out of his tractor to take a telephone call, to attend a meeting, or to write a memo. *What* is to be done is always obvious in manual work.

However, in knowledge work the task does not program the worker. A major crisis in a hospital, such as when a patient suddenly goes into coma, does of course control the nurse's task and programs her; but otherwise, it is largely the nurse's decision whether to spend time at the patient bed or whether to spend time filling out papers. Engineers are constantly being pulled off their task by having to write a report or rewrite it, by being asked to attend a meeting, and so on. The job of the salesperson in the department store is to serve the customer and to provide the merchandise the customer is interested in or should become interested in. Instead, the salesperson spends an enormous amount of time on paperwork, on checking whether merchandise is in stock, on checking when and how it can be delivered, and so on—all things that take salespeople away from the customer and do not add anything to their productivity in doing what salespeople are being paid for, which is to sell and to satisfy the customer.

The first requirement in tackling knowledge work is to find out what the task is so as to make it possible to concentrate knowledge workers on the task and to eliminate everything else—at least as far as it can possibly be eliminated. This requires that the knowledge workers themselves define what the task is or should be—and only the knowledge workers themselves can do that. Work on knowledge-worker productivity therefore begins with asking the knowledge workers themselves: *What is your task? What should it be? What should you be expected to contribute?* and *What hampers you in doing your task and should be eliminated?*

Knowledge workers themselves almost always have thought through these questions and can answer them. Still, it then usually takes time and hard work to restructure their jobs so that they can actually make the contribution they are already being paid for. However, asking the questions and taking action on the answers usually doubles or triples knowledge-worker productivity, and quite fast.

Nurses in a major hospital were asked these questions. They were sharply divided as to what their task was, with one group saying "patient care" and another saying "satisfying the physicians." However, they were in complete agreement on the things that made them unproductive. They called them "chores"—paperwork, arranging flowers, answering the phone calls of patients' relatives, answering the patients' bells, and so on. All—or nearly all—of these could be

turned over to a non-nurse floor clerk, paid a fraction of a nurse's pay. The productivity of the nurses on the floor immediately more than doubled, as measured by the time nurses spent at the patients' beds. Patient satisfaction more than doubled and turnover of nurses (which had been catastrophically high) almost disappeared—all within four months.

Once the task has been defined, the next requirements can be tackled, and they will be tackled by the knowledge workers themselves. These requirements are:

- Knowledge workers' *responsibility* for their own contribution. It is the knowledge worker's decision what he or she should be held accountable for in terms of quality and quantity with respect to time and with respect to cost. Knowledge workers have to have autonomy and that entails responsibility.
- Continuous innovation *has to be built into the knowledge worker's job*.
- *Continuous learning* and *continuous teaching* have to be built into the job.

One central requirement of knowledge-worker productivity remains. We have to answer the question: What is quality? In some knowledge work—and especially in some work requiring a high degree of knowledge—we already measure quality. Surgeons, for example, are routinely measured, especially by their colleagues, by their success rates in difficult and dangerous procedures (e.g., by the survival rates of their open-heart surgical patients or the full recovery rates of their orthopedic-surgery patients). By and large, we mainly have judgments rather than measures regarding the quality of a great deal of knowledge work. The main trouble is, however, not the difficulty of measuring quality. It is the difficulty—and more particularly the sharp disagreements—in defining what the task is and what it should be.

The best example of this is the American school system. As every one knows, public schools in the American inner city have become disaster areas. Next to them—in the same location and serving the same kind of children—are private (mostly Christian) schools in which the kids behave well and learn well. There is endless speculation to explain these enormous quality differences. A major reason is surely that the two kinds of school define their tasks differently. The typical public school defines its task as “helping the underprivileged,” while the typical private school (and especially the Parochial Schools of the Catholic Church) define their task as “enabling those who want to learn, to learn.” One therefore is governed by its scholastic failures, the other one by its scholastic successes.

Similarly, the research departments at two major pharmaceutical companies have totally different results because they define their tasks differently. One sees its task as not having failures, that is, in working steadily on fairly minor but predictable improvements in existing products and for established markets. The other one defines its task as producing “breakthroughs” and therefore courts risks. Both are considered fairly successful—by themselves, by their own top

managements, and by outside analysts. Yet each operates quite differently and quite differently defines its own productivity and that of its research scientists.

To define quality in knowledge work and to convert the definition into knowledge-worker productivity is thus to a large extent a matter of defining the task. It requires the difficult, risk-taking, and always controversial definition as to what “results” are for a given enterprise and a given activity. We therefore actually *know* how to do it. Nevertheless, the question is a completely new one for most organizations and also for most knowledge workers. To answer it *requires* controversy, *requires* dissent.

The Knowledge Worker as Capital Asset

In no other area is the difference greater between manual-worker productivity and knowledge-worker productivity than in their respective *economics*. Economic theory and most business practice sees manual workers as a *cost*. To be productive, knowledge workers must be considered a *capital asset*. Costs need to be controlled and reduced. Assets need to be made to grow.

To be sure, in managing manual workers we learned fairly early that high turnover (i.e., losing workers) is very costly. The Ford Motor Company, as is well known, increased the pay of skilled workers from eighty cents a day to \$5.00 a day on January 1, 1914. It did so because its turnover had been so excessive as to make its labor costs prohibitively high; it had to hire 60,000 people a year to keep 10,000. Even so, everybody (including Henry Ford himself, who had at first been bitterly opposed to this increase) was convinced that the higher wages would greatly reduce the company’s profits. Instead, in the very first year, profits almost doubled. Paid \$5.00 a day, practically no workers left—in fact, the Ford Motor Company soon had a waiting list.

However, short of the costs of turnover, rehiring, retraining, and so on, the manual worker is still being seen as a cost. This is true even in Japan, despite the emphasis on lifetime employment and on building a “loyal,” permanent work force. The management of people at work, based on millennia of work being almost totally manual work, still assumes that with few exceptions (e.g., highly skilled people) one manual worker is like any other manual worker.

This is definitely not true for knowledge work. Employees who do manual work do not own the means of production. They may, and often do, have a lot of valuable experience, but that experience is valuable only at the place where they work. It is not portable. Knowledge workers, however, *own* the means of production. That knowledge between their ears is a totally portable and enormous capital asset. Because knowledge workers own their means of production, they are mobile. It may not be true for most of them that the organization needs them more than they need the organization. For most of them it is a symbiotic relationship in which they need each other in equal measure. It is

not true, as it was for the manual worker in modern industry, that they need the job much more than the job needs them.

Management's job is to preserve the assets of the institution in its care. What does this mean when the knowledge of the individual knowledge worker becomes an asset—and, in more and more cases, the *main* asset—of an institution? What does this mean for personnel policy? What is needed to attract and to hold the highest producing knowledge workers? What is needed to increase their productivity and to convert their increased productivity into performance capacity for the organization?

The Technologists

A very large number of knowledge workers do both knowledge work *and* manual work. I call them “technologists.” This group includes people who apply knowledge of the highest order.

Surgeons preparing for an operation to correct a brain aneurysm before it produces a lethal brain hemorrhage, spend hours in diagnosis *before* they cut—and that requires specialized knowledge of the highest order. Again, during the surgery, an unexpected complication may occur which calls for theoretical knowledge and judgment, both of the very highest order. However, the surgery itself is manual work—and manual work consisting of repetitive, manual operations in which the emphasis is on speed, accuracy, and uniformity. These operations are studied, organized, learned, and practiced exactly like any manual work—that is, by the same methods Taylor first developed for factory work.

The technologist group also contains large numbers of people in whose work knowledge is relatively subordinate—though it is always crucial. The file clerk's job—and that of the clerk's computer-operator successor—requires a knowledge of the alphabet that no experience can teach. This knowledge is a small part of an otherwise manual task, but it is its foundation and is absolutely crucial.

Technologists may be the single biggest group of knowledge workers. They may also be the fastest-growing group. They include the great majority of health-care workers: lab-technicians; rehabilitation technicians; technicians in imaging such as X-ray, ultrasound, magnetic-resonance imaging; and so on. They include dentists and all dental-support people. They include automobile mechanics and all kinds of repair and installation people. In fact, the technologist may be the true successor to the 19th and 20th century skilled workers.

Technologists are also the one group in which developed countries can have a true and long-lasting competitive advantage. When it comes to truly high knowledge, no country can any longer have much of a lead the way 19th century Germany had through its University. Among theoretical physicists, mathematicians, economic theorists, and the like, there is no “nationality.” Any country can, at fairly low cost, train a substantial number of high-knowledge

people. India, for instance, despite her poverty, has been training fairly large numbers of first-rate physicians and first-rate computer programmers. Similarly, there is no “nationality” in respect to the productivity of manual labor. Training based on Scientific Management has made all countries capable of attaining—overnight—the manual-worker productivity of the most advanced country, industry, or company. Only by educating technologists can the developed countries still have a meaningful and lasting competitive edge.

The U.S. is so far the only country that has developed this advantage through its unique nationwide systems of community colleges. The community college was actually *designed* (beginning in the 1920s) to educate technologists who have *both* the needed theoretical knowledge *and* the manual skill. On this, I am convinced, rests both the still huge productivity advantage of the American economy and the (so far unique) American ability to create, almost overnight, new and different industries.

Currently, nothing quite like the American Community College exists in any other nation. The famous Japanese school system produces either people prepared only for manual work or people prepared only for knowledge work. Not until the year 2003 is the first Japanese institution devoted to train technologists supposed to get started. The even more famous German apprenticeship system (started in the 1830s) was one of the main factors in Germany’s becoming the world’s leading manufacturer. However, it focused—and still focuses—primarily on manual skills and slight theoretical knowledge. It is thus in danger of becoming rapidly obsolete.

Other developed countries should be expected to catch up with the U.S. fairly fast. “Emerging” or “Third World” countries are, however, likely to be decades behind—in part because educating technologists is expensive, in part because in these countries people of knowledge still look down with disdain, if not with contempt, on working with one’s hands. “That’s what we have servants for” is still their prevailing attitude. However, in developed countries—and again foremost in the U.S.—more and more manual workers are going to be technologists. To increase knowledge-worker productivity, increasing the productivity of technologists deserves to be given high priority.

The job was actually done more than seventy years ago by the American Telephone Company (AT&T) for its technologists, the people who install, maintain, and replace telephones. By the early 1920s, the technologists working outside the telephone office and at the customer’s location had become a major cost center—and at the same time a major cause of customer unhappiness and dissatisfaction. It took about five years or so (from 1920 until 1925) for AT&T—which had by that time acquired a near monopoly on providing telephone service in the United States and in parts of Canada—to realize that the task was not installing, maintaining, repairing, and replacing telephones and telephone connections. *The task was to create a satisfied customer.* Once they realized this, it became fairly easy to organize the job. It meant, first, that the technicians themselves

had to define what “satisfaction” meant. The results were standards that established that every order for a new telephone or an additional telephone connection would have to be satisfied within 48 hours, and that every request for repair would have to be satisfied the same day if made before noon, or by noon the following day. Then it became clear that the individual service people—in those days all men, of course—would have to be active participants in such decisions as whether to have one person installing and replacing telephones and another one maintaining and repairing them or whether the same people had to be able to do all jobs (which in the end turned out to be the right answer). Then these people had to be taught a very substantial amount of theoretical knowledge—and in those days few of them had more than six years of schooling. They had to understand how a telephone works. They had to understand how a switchboard works. They had to understand how the telephone system works. These people were not qualified engineers nor skilled craftsmen, but they had to know enough electronics to diagnose unexpected problems and be able to cope with them. Then they were trained in the repetitive manual operation or in the “one right way” (that is, through the methods of Scientific Management) and *they* made the decisions (e.g., where and how to connect the individual telephone to the system and what particular kind of telephone and service would be the most suitable for a given home or a given office). They had to become salesmen in addition to being servicemen.

Finally, the telephone company faced the problem how to define *quality*. The technologist had to work by himself. He could not be supervised. He, therefore, had to define quality, and he had to deliver it. It took another several years before that was answered. At first the telephone company thought that this meant a sample check, which had supervisors go out and look at a sample (maybe every 20th or 30th job done by an individual service person) and check it for quality. This very soon turned out to be the wrong way of doing the job, annoying servicemen and customers alike. Then the telephone company defined quality as “no complaints”—and they soon found out that only extremely unhappy customers complained. It then had to redefine quality as “positive customer satisfaction.” In the end, this then meant that the serviceman himself controlled quality (e.g., by calling up a week or ten days after he had done a job and asking the customer whether the work was satisfactory and whether there was anything more the technician could possibly do to give the customer the best possible and most satisfactory service).

I have intentionally gone into considerable detail in describing this early example because it exemplifies the three elements for making the worker who is both a knowledge worker and a manual worker both effective and productive.

- First, there is the answer to the question “What is the task?”—the key question in making every knowledge worker more productive. As the example of the Bell System shows, this is not an obvious answer. As the Bell System people learned, the only people who knew the answer to this were the technologists themselves. In fact, until they asked the

technologists, they floundered. However, as soon as the technologists were asked, the answer came back loud and clear: “a satisfied customer.”

- Then, the technologists had to take full responsibility for giving customer satisfaction, that is, for delivering quality. This showed what *formal knowledge* the technologist needed. Only then could the *manual* part of the job be organized for manual-worker productivity.
- Above all, this example shows that technologists have to be treated as *knowledge workers*. No matter how important the manual part of their work—and it may take as much time as it did in the case of the AT&T installers—the focus has to be on making the technologist knowledgeable, responsible, and productive as a knowledge worker.

Knowledge Work as a System

Productivity of the knowledge worker will almost always require that the *work itself* be restructured and be made part of a *system*. One example is servicing expensive equipment, such as huge and expensive earth-moving machines. Traditionally, this had been seen as distinct and separate from the job of making and selling the machines. However, when the U.S. Caterpillar Company, the world’s largest producer of such equipment, asked “What are we getting paid for?” the answer was “We are not getting paid for machinery. We are getting paid for what the machinery does at the customer’s place of business. That means keeping the equipment running, since even one hour during which the equipment is out of operation may cost the customer far more than the equipment itself.” In other words, the answer to “What is our business?” was “Service.” This then led to a total restructuring of operations all the way back to the factory in order that the customer could be guaranteed continuing operations and immediate repairs or replacements. The service representative, usually a technologist, has become the true “decision maker.”

As another example, a group of about 25 orthopedic surgeons in a Midwestern U.S. city, have organized themselves as a “system” to: produce the highest quality work; make optimal use of the limited and expensive resources of operating and recovery rooms; make optimal use of the supporting knowledge people such as anesthesiologists or surgical nurses; build continuous learning and continuous innovation into the work of the entire group and of every member thereof; and, finally, minimize costs. Each of the surgeons retains full control of his or her practice. He or she is fully responsible for obtaining and treating the individual patient. Traditionally, surgeons schedule surgeries early in the morning. Hence, operating rooms and recovery rooms are standing empty most of the time. The group now schedules the use of operating and recovery rooms for the entire group so that this scarce and extremely expensive resource is utilized ten hours a day. The group, as a group, decides on the standardization of tools and equipment so as to obtain the highest quality at the lowest cost. Finally, the group has also built quality control into its system. Every three months three

surgeons are designated to scrutinize every operation done by each of the members—the diagnosis, the surgery, the after-treatment. They then sit down with the individual surgeons and discuss their performance. They suggest where there is need for improvement and they also may recommend that a certain surgeon be asked to leave the group when his or her work is not satisfactory. Each year, the quality standards that these supervising committees apply are discussed with the whole group and are raised, often substantially. As a result, this group now does almost four times as much work as it did before. It has cut the costs by 50%, half of it by cutting back on the waste of operating and recovery rooms and half by standardizing tools and equipment. In such measurable areas as success rates in knee or shoulder replacements and in recovery after sports injuries, it has greatly improved its results.

What to do about knowledge-worker productivity is thus largely known. So is *how* to do it.

How to Begin?

Making knowledge workers more productive requires changes in basic attitude, whereas making the manual worker more productive only required telling the worker how to do the job. Furthermore, making knowledge workers more productive requires changes in attitude not only on the part of the individual knowledge worker, but on the part of the whole organization. It therefore has to be “piloted,” as any major change should be. The first step is to find an area in the organization where there is a group of knowledge workers who are receptive. (The orthopedic surgeons, for instance, first had their new ideas tried out by four physicians who had long argued for radical changes.) The next step is to work consistently, patiently, and for a considerable length of time with this small group. The first attempts, even if greeted with great enthusiasm, will almost certainly run into all kinds of unexpected problems. It is only after the productivity of this small group of knowledge workers has been substantially increased that the new ways of doing the work can be extended to a larger area, if not to the entire organization. At this point, the main problems will be known, such as where resistance can be expected (e.g., from middle management) or what changes in task, organization, measurements, or attitudes are needed for full effectiveness. To bypass the pilot stage—and there is always pressure to do so—only means that the mistakes become public while the successes stay hidden. It only means discrediting the entire enterprise. If properly piloted, a great deal can be done to improve knowledge-worker productivity.

Knowledge-worker productivity is the biggest of the 21st-century management challenges. In the developed countries, it is their first *survival requirement*. In no other way can the developed countries hope to maintain themselves, let alone maintain their leadership and their standards of living. In the 20th century, this leadership very largely depended on making the manual worker more productive. Any country, any industry, any business can do that today using the

methods that the developed countries have worked out and put into practice in the 120 years since Frederick Winslow Taylor first looked at manual work. Anybody today, any place, can apply those policies to training, the organization of work, and the productivity of workers—even if they are barely literate, if not illiterate, and totally unskilled.

Above all, the supply of young people available for manual work will be rapidly shrinking in the developed countries—in the West and in Japan very fast, in the U.S. somewhat more slowly—whereas the supply of such people will still grow fast in the emerging and developing countries for at least another thirty or forty years. The only possible advantage developed countries can hope to have is in the supply of people prepared, educated, and trained for knowledge work. There, for another fifty years, the developed countries can expect to have substantial advantages, both in quality and in quantity. Whether this advantage will translate into performance depends on the ability of the developed countries—and of every industry in it, of every company in it, of every institution in it—to raise the productivity of the knowledge worker and to raise it as fast as the developed countries have raised the productivity of the manual worker in the last hundred years.

The countries and the industries that have emerged as the leaders in the last hundred years in the world are the countries and the industries that have led in raising the productivity of the manual worker—the U.S. first, Japan and Germany second. Fifty years from now, if not much sooner, 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.

The Governance of the Corporation

What does the emergence of the knowledge worker and of knowledge-worker productivity mean for the *governance of the corporation*? What does it mean for the future and structure of the economic system?

In the last ten or fifteen years, pension funds and other institutional investors became the main share owners of the equity capital of publicly owned companies in all developed countries. In the U.S., this has triggered a furious debate on the governance of corporations. With the emergence of pension funds and mutual funds as the owners of publicly owned companies, power has shifted to these new owners. Similar shifts in both the definition of the purpose of economic organizations (such as the business corporation) and their governance can be expected to occur in all developed countries.

Within a fairly short period of time, we will face the problem of the governance of corporations again. We will have to redefine the purpose of the employing organization and of its management as *both* satisfying the legal owners (such as shareholders) and satisfying the owners of the human capital that gives the organization its wealth-producing power—that is, satisfying the knowledge

workers. Increasingly, the ability of organizations—and not only of businesses—to survive will come to depend on their “comparative advantage” in making the knowledge worker more productive. The ability to attract and hold the best of the knowledge workers is the first and most fundamental precondition.

However, can this be *measured* or is it purely an “intangible”? This will surely be a central problem for management, for investors, and for capital markets. What does “capitalism” mean when knowledge governs rather than money? And what do “free markets” mean when knowledge workers—and no one else can “own” knowledge—are the true assets? Knowledge workers can neither be bought nor be sold. They do not come with a merger or an acquisition. In fact, although they are the greatest “value,” they have no “market value”—that means, of course, that they are not an “asset” in any sense of the term.

These questions go far beyond the scope of this article. However, it is certain that the emergence as *key questions* of the knowledge worker and of the knowledge-worker’s productivity will, within a few decades, bring about fundamental changes in the very structure and nature of the economic system.

Notes

1. For work in the oldest knowledge profession—that is, in Medicine—Taylor’s close contemporary William Osler (1849-1919) did what Taylor did and at the same time in his 1892 book *The Principles and Practice of Medicine* (arguably the best textbook since Euclid’s *Geometry* in the third century B.C.). Osler’s work has rightly been called the application of Scientific Management to Medical Diagnosis. Like Taylor, Osler preached that there is no “skill,” there is only *method*.