

Technology, Work, and the Organization: The Impact of Expert Systems

Rob R. Weitz

"Over the last decade a new technology has begun to take hold in . . . business, one so new that its significance is still difficult to evaluate. While many aspects of this technology are uncertain, it seems clear that it will move into the managerial scene rapidly, with definite and far reaching impact on managerial organization."

—Harold Leavitt
and Thomas
Whisler
"Management in the 1980s"
Harvard Business Review
November–December 1958

This article examines the near-term impact of expert system technology on work and the organization. First, an approach is taken for forecasting the likely extent of the diffusion, or success, of the technology. Next, the case of advanced manufacturing technologies and their effects is considered. From this analysis, a framework is constructed for viewing the impact of these technologies—and technologies in general—as a function of the technology itself; market realities; and personal, organizational, and societal values and policy choices. Two scenarios are proposed with respect to the application of this framework to expert systems. The first concludes that expert systems will have little impact on the nature of work and the organization. The second scenario posits that expert system diffusion will be pulled by, and will be a contributing factor toward, the evolution of the lean, flexible, knowledge-intensive, postindustrial organization.

Some 30 years after the publication of the article that began with this quotation, the impact of computer-based technology on business is still being discussed, debated, predicted, and assessed. Indeed, although the technology itself advances at seemingly breakneck speed, the questions surrounding the technology have remained remarkably stagnant: How will (or should) organizational structures change? How will current jobs and tasks change? Will some jobs disappear altogether? Which jobs will be cre-



The focus . . . is on . . . the impact of one technology—expert systems—on nonskilled, semiskilled, and skilled workers and first-line management over the coming three to five years.

ated? What net change in the number of jobs will occur? What can be expected as the ripples of secondary and tertiary effects—social, economic, and political—are realized? How can government, industry, academia, and labor best plan for (or direct) technological innovations and their impact?

These questions constitute an imposing intellectual pie. The focus of this article is on but a small piece of this pie: the impact of one technology—expert systems—on nonskilled, semiskilled, and skilled workers and first-line management over the coming three to five years.

Assessing trends and future events is a forecasting problem. This article approaches estimating the impact of expert systems from this perspective. Two different but related forecasts are required: The first is an estimate of the extent of diffusion or success of expert system technology in the marketplace, and the second involves determining the actual impact of the technology once in place. Although I do tread on dangerous ground and suggest possible future scenarios, the major thrust of this work is to stipulate an appropriate methodology for forecasting the impact of a technology in general and expert systems in particular.

A large body of research exists that focuses on information technology and organizational change. In surveying the literature, Markus and Robey (1988) note the difficulties in drawing reliable generalizations from this research. Here, I argue that the diffusion and impact of a technology is never purely a function of the technology itself. Rather, the extent and manner of use of a technology is a complex function of the technology, market forces, and the values and policy choices of the society, organization, and individuals that make up the environment in which the technology is introduced.

This article is organized in the following fashion: The first section, Expert Systems, establishes expert systems as a definable technology with enumerable benefits and limitations. Section 2, Forecasting, discusses the artful science of forecasting and begins to specify the implications for this work. The third section, Forecasting New Markets, con-

tinues this theme, specifically looking at the forecasting process for new products and markets. Section 4, Advanced Manufacturing Technology, outlines the current manufacturing environment and attempts to draw parallels between recent technological advances in manufacturing and expert systems. A model for better understanding and communicating the impact of a technology is proposed in section 5, Implications; then, the focus is drawn back to expert systems with the suggestion of two likely scenarios over the three- to five-year horizon. The final section, Observations and Questions, provides a summary and suggests directions for future research.

Expert Systems

"I have one final thing to tell you; I am not actually a human being. I am in fact a robot built by Marvin Minsky in 1968."

—Michael Arbib

Professor of Biomedical Engineering,
Computer Science, Neurobiology, and
Physiology

At a conference on the mind-body problem, 1987

The terms knowledge based and expert system denote a particular technology. (For defining criteria, see Harmon and King [1985] or Hayes-Roth, Waterman, and Lenat [1983].) As a technology, expert systems provide certain benefits at certain costs and within certain limitations. The success of a technology depends at least in part on these factors. It is instructive to recall both the benefits and limitations of expert systems. On the positive side, expert systems (1) can solve or assist in solving difficult problems that have eluded previous solution by other computer-based technologies; (2) can perform with high reliability and consistency; (3) allow for experts to concentrate on rarer, more interesting tasks; (4) provide a central, standard, accessible cache of knowledge; (5) make knowledge explicit, thereby promoting knowledge enhancement; and (6) can provide a tool for training neophytes.

On the negative side, expert systems (1) work in relatively narrow fields of expertise (they are not general problem solvers); (2) are



Forecasting requires predicting what will happen and when it will happen along with an associated statement of confidence.

suitable for problems that are difficult but not too difficult; (3) rely for their construction on expertise being recognized as both useful and available for the problem selected; (4) are resource (that is, time, money, and people) intensive to build; (5) generally do not learn; (6) do not know the limits of their own expertise; and (7) reason from superficial or surface knowledge.

Forecasting

"I think there is a world market for about five computers."

—Thomas J. Watson, Sr.

Chairman of the Board of IBM, 1943

Forecasts such as the one expressed in the previous quotation, absurd in retrospect, are easy to find. Many events have happened that were not predicted significantly in advance of their occurrence (the popularity of jogging shoes, the number of women in the U.S. work force, the Walkman), and many predictions simply never materialized (free, unlimited energy; automated computer translation; the running out of oil). Further, expertise in the domain in question does not correlate well with better forecasting. Witness Henry Warner's comment in 1921: "Who the hell wants to hear actors talk?" See Makridakis (1988) for these and other examples. Why is this the case? The evidence points to human information-processing limitations and biases. Working from the psychological literature on human judgmental abilities, Hogarth and Makridakis (1981) cite, among a litany of other factors, "the 'illusion of control', accumulation of redundant information, failure to seek possible disconfirming evidence, and overconfidence in judgement (as) liable to induce serious errors in forecasting and planning" (p. 115).

Forecasting requires predicting what will happen and when it will happen along with an associated statement of confidence. Claiming that due to expert systems, there is a 90-percent chance that first-line managers will disappear within the next three to five years

is an example. For certain regular events, such as predicting when the sun will rise tomorrow, forecasting is easy. For less regular events and longer time horizons, the task is more difficult. A large body of empirical and theoretical research in the field of forecasting specifies the factors in a particular problem contributing to forecast difficulty and, given the difficulties, how one should proceed to forecast. See Armstrong (1985) and Makridakis (1986), for example.

Certain factors contribute to the difficulty of the problem at hand; these factors are current status, forecasting horizon, and methods.

Current Status

To forecast where we will be in the future (say, in three to five years), it is always helpful to know where we are now. In this case, we don't. That is to say, the extent to which verifiable expert system technology has permeated the workplace is at the moment less than well known. A pervasive argument exists that few working expert systems are publicized due to the strategic value of this resource to the company. (That is, much as with male bravado, those who do it don't talk about it, and those who talk about it don't do it.) The counterargument is that expert systems are difficult and expensive to develop though limited in application and, therefore, will not be widely used until the underlying technology takes a quantum leap forward. Further, given the relatively small sample of verifiable, successful applications of expert system technology, it is naturally difficult to generalize about the currently observed impact of expert system introduction.

Forecasting Horizon

Forecasting horizon is inversely related to forecasting accuracy. This result is intuitive: The likelihood of unforeseen events increases as the time period over which the forecast is made increases. In this case, three to five years is, generally speaking, neither short term nor long term but, rather, medium term. This three- to five-year window implies that the likelihood of the diffusion of some technological breakthrough, as described in the previous subsection, will be low over the forecasting horizon. A longer time horizon would force a methodology that allowed for technological forecasting. Under the assumed conditions, some form of an extrapolation of recent trends is possible.

In one way, the long-term forecast is easier for this problem. The trend is for expert sys-

tems to substitute or supplement human mental tasks. The hard part is determining the rate at which this trend will happen.

A final note is in order here: For the medium term time horizon, people tend to overestimate the amount of change caused by technology. Several phenomena can account for this overestimation. First, optimism, wishful thinking, and a desire to control the world in which we live tend to make people believe in the power of technology and mankind's abilities (through technology) to mold a better future. Second, press reports, with the intent of making interesting reading, tend to exaggerate the capabilities and potential applications of new technologies and minimize limitations, the required time to commercialization, and the potential problems involved. Third, typically, it's the successful technological breakthroughs that stand out in one's memory and the unsuccessful or less successful ones that tend to fade with time. New technologies tend to be associated with those previous breakthroughs that can be recalled. Ultimately, only in the long term can the actual success and impact of a technology be accurately observed.

Methods

Qualitative forecasting, that is, forecasting based on human judgment, is difficult, as described previously, due to human biases and the limitations of human information processing. To illustrate the point, when quantitative methods are feasible, for example, in forecasting time series or when using econometric techniques, the empirical evidence indicates that (human) judgmental methods are not more accurate than even the simplest of these quantitative methods. Nevertheless, there are problems such as this one where there is no recourse but to rely on judgment. (What is needed is an expert system for assisting in the forecasting process! For a view of one such system, see Weitz [1986].) Makridakis, Wheelwright, and McGee (1983) describe a number of these methods, among them the jury of executive opinion, surveys, scenario development, and analogy. A jury in the qualitative forecasting sense is simply a group of knowledgeable individuals gathered in one place that through their interaction come to some estimate of future outcomes and their likely occurrence. In scenario development, multiple sets of plausible assumptions are made, and the likely outcomes for each set of assumptions are derived. Whether scenario development has predictive validity, it is certainly useful for

contingency planning. As the name suggests, analogy methods draw on previously occurring circumstances similar to the one in question. The sequence of events in the analogy is seen as a guide to likely future patterns. For example, studies of the life cycle of many products indicate an S-shaped growth pattern. In forecasting the sales of a relatively new product, this rate of growth analogy can provide guidance.

None of these methods is known for its stunning successes. The growth rate analogy example points to some of the difficulties. One must assume the new product will exemplify the growth characteristics of this shape and then predict the size of the S. Before release—or even early in the product release—this determination is not trivial.

Forecasting New Markets

"There doesn't seem to be any real limit to the growth of the computer industry."

—Thomas J. Watson, Jr.

Chairman of the Board of IBM, 1968

To reiterate, the problem here is to forecast the diffusion of the technology and its impact; the two are different but not independent. Widespread effects would be impossible without large-scale exploitation of expert systems. However, it is conceivable, and therefore, the possibility should be considered, that wide acceptance and use of expert systems could simply result in small, incremental changes in the nature of work and the organization—that these systems, for example, will simply aid and enhance current tasks and processes. I presume, therefore, that significant exploitation of expert systems is a necessary but not sufficient condition for widespread impact on work and the organization. It should be clear, however, that moderate or small-scale utilization of expert system technology could have significant, although localized, effects (that is, within certain companies or industries) or not very significant effects.

In the previous section, I outlined issues of importance in forecasting and how they apply to this problem in a general sense. In this section, I focus on the particular guidelines involved in forecasting the success of a technology or product. The level of impact is addressed later in this article.

Will the diffusion of expert system technology be widespread? Forecasting the future extent of use can be viewed as trying to determine if expert systems are a strong growth market. Schnaars and Berenson (1986) systematically sampled the popular business press between 1960 and 1979 and assessed the

*... optimism,
wishful
thinking, and
a desire to
control the
world in
which we live
tend to make
people believe
in the power
of technology
and
mankind's
abilities . . .
to mold a
better future.*

Forecast	Year Made	Time Horizon (in years)
Tooth Decay Vaccine	1968	10
Rotary Authmobile Engines	1971	4
Ultrasonic Dishwashers	1967	5
Turning fuel (i.e., oil) into food	1961	Unspecified

Figure 1. Some Mistaken Growth Market Forecasts (Schnaars and Berenson 1986).

growth market forecasts appearing there. The purpose was to see which lessons, if any, could be drawn from the successes and failures they encountered. The study included forecasts for “significant new products, markets, and emerging technologies.” Of the 90 growth markets in the sample, 48 never materialized, and 42 were successful forecasts. Figure 1 highlights some of these mistaken forecasts.

The primary reasons cited for the failure of these growth market forecasts are (1) an overvaluation of technology, that is, an infatuation with technology for technology’s sake; (2) a failure to realize that the product or technology did not offer a real benefit to the consumer at a reasonable cost; and (3) an inability to predict social or demographic changes that significantly altered the relative advantage (for example, the cost-benefit) of the product or technology to the consumer.

Some of the successful forecasts are presented in figure 2. Correct forecasts tended to be demographic in nature and those where fundamental market research was undertaken and where a real benefit was offered to the consumer at the price offered.

It was relatively easy to predict the number of senior citizens or baby boom adults, for example, by simple extrapolation techniques. The obvious needs of these groups could then be predicted without great difficulty. Where clear demands of cohort groups could not be determined, good market research pointed the way to possible changes in historical patterns of consumer behavior. Finally, successful markets involved providing the consumer with a real cost-benefit advantage.

The authors conclude with the following guidelines for successfully predicting growth markets: (1) stress market fundamentals; (2) be aware of the assumptions underlying the forecast, particularly those based on immutable trends; (3) avoid being dazzled by

Forecasts
The Elderly Market
The Baby Boom Market
Personal Computers
Microwave Ovens
Home Pregnancy Tests

Figure 2. Some Correct Growth Market Forecasts (Schnaars and Berenson 1986).

technology; and (4) avoid forecasts made by those “in bed” with the technology. Finally, the use of scenario analysis might be helpful.

The study emphasizes that

regardless of technological dazzle, true growth markets have been driven by fundamentals: Who are the customers? How large is the market? Will the proffered technology offer them a real benefit over existing and subsequent substitutes? Is the derived benefit worth the price you will have to charge? Are cost efficiencies probable? Are social trends moving towards or away from this product? Place the bulk of your effort into answering these questions. The potential for the market depends more on these issues than any other factor. (p. 87)

In the next section, I draw on the manufacturing environment to illustrate the interplay of new technologies, value systems, economic realities, and the nature of work and the organization. In focusing on the current experience with new technologies in manufacturing, I attempt to draw an analogy for the extent and impact of expert system introduction in the workplace. It should be kept in mind that manufacturing includes a host of tasks ranging from the nonskilled to the very skilled and up through management. Skilled tasks include such areas as design, scheduling, monitoring, maintenance, and repair—ideal domains for expert systems.

Advanced Manufacturing Technology

“What is the sound of one hand?”
—Zen Riddle

Attributed to Hakuin Ekaku (1686–1769)
What is it about the Japanese that has made their manufacturing so successful? The

answer to this more recent riddle is more forthcoming and can be summarized under the headings technology, just-in-time (JIT) scheduling policies, and participatory management. The technology includes robotics and computer-aided design, manufacture, and process planning. The technology itself, however, does not account for the success of the Japanese. Linked to the technology is a set of policies that allow for its maximal utilization.

JIT scheduling revokes the decades-old notion that maintaining sizable inventories (of cars at dealerships, for example) is necessary to smooth response to demand. Although the traditional mechanism for reducing carrying costs was to minimize inventory, JIT dictates the ultimate minimum in-process inventory level: zero. The implications are enormous: Rapid, flexible manufacturing processes, from supplier to customer, must be assured.

Participatory management breaks from Taylor's *scientific management*, developed at the beginning of the century. Taylorism is perhaps best exemplified by Henry Ford's assembly line: Intellectual content is removed from work, required movements of the worker are minimized and regularized, and a distinct separation between hourly workers and their salaried supervisors is enforced. Participatory schemes, although taking different forms, typically have workers doing a variety of tasks, with some intellectual satisfaction and sense of accomplishment associated with the tasks. For example, quality control (which under a Taylorized system would be performed after assembly by supervisors) is done by the workers themselves as part of their responsibilities during assembly. In Japan, a chief component of this management policy is that workers are given the time, incentive, and resources to reflect on and improve their jobs. Hayes and Wheelwright (1984) treat these topics in depth.

It has been argued that the Japanese developed these policies as an outgrowth of the values and ethics inherent in their society. These underlie a philosophy that, for example, promotes a team spirit and a sense of orderliness. An important question then is to what extent are the policies directing the acceptance and use of a technology inextricably linked to the technology itself and to what extent are the policies an outgrowth of the country-company culture in which they originate?

Walton and Susman (1987) report on a study of 24 plants that are pioneering the use of advanced manufacturing technologies (AMTs) in the United States. The focus of the

In Japan . . . workers are given the time, incentive, and resources to reflect on and improve their jobs.

report is on human resource strategies associated with the new technologies. When AMTs are introduced in the workplace, they observe

Closer interdependence of activities.

Different skill requirements—usually higher average skill levels.

More immediate—and more costly—consequences of any malfunction.

Output more sensitive to variations in human skills, knowledge, and attitudes, and to mental effort rather than physical effort.

More dynamism, that is, continual change and development.

Higher capital investment per employee and fewer employees responsible for a particular product, part, or process.

Although AMTs have generally been introduced into a company on a piecemeal basis, the growth and interdependence of automation and information technology make essential

A highly skilled, flexible, coordinated, and committed workforce.

Lean, flat, flexible, and innovative management.

The ability to retain developed talent.

A strong partnership between management and labor unions. (p. 98)

The technology, coupled with economic imperatives, has led these companies to reduce the number of job classifications while broadening them in scope, promote team structures, revamp compensation and appraisal systems, and revise (or consider revising) training and selection procedures. JIT and participatory management are, of course, implicit in this package.

A general upgrading of skill requirements is observed with operators handling multiple functions, some of them previously performed by technicians; the technicians, in turn, took on responsibilities previously afforded only to engineers. The task of supervisor has evolved into one of facilitator as opposed to task master. Lower management as well has seen its responsibilities upgraded,




Figure 3. Impact of Technology Matrix.

broadened, or both. The authors note, however, that

the policy choices. . . . constitute a strategy to upgrade workers' knowledge and skills. Had management simply followed practice, the more demanding AMT tasks in most of the companies we reviewed probably would have been assigned to managers or other professionals. Whether they use the term explicitly or not, companies seem to be weighing benefits and risks and deciding in favor of upgrading. Downgrading, however, remains a live option in the minds of some managements, and the choice is still a judgement call (emphasis mine). (pp. 103–104)

The authors note that the new technologies and practices require fewer overall jobs. Because flexibility and the ability to learn are fundamentally important qualifications, the implication is that workers without these abilities lose out in the new manufacturing world.

In summary, although the technology and today's business environment seem to force certain policies, others apparently remain a function of choice.

Implications

"... we predict that [in the 1980s] large industrial organizations will recentralize. . . ."

—Harold Leavitt and Thomas Whisler
"Management in the 1980s"
Harvard Business Review

November–December 1958

"If you look at the way truly decentralized companies outperform centralized compa-

nies, you've got to say it's a trend that's got to force itself on even the most reluctant CEOs."

—Thomas J. Peters,
Coauthor of *In Search of Excellence*
Business Week
April 25, 1983

As Peters implies, and others in the popular press echo, the successful organization today, whether manufacturing, is lean, flexible, and (correspondingly) decentralized. Huber (1984) refers to the

post-industrial society [that] will be characterized by more and increasing knowledge, more and increasing complexity, and more and increasing turbulence. These, in combination, will pose an organizational environment qualitatively more demanding than those in our experience. (p. 931)

This is the overall business environment in which organizations currently find themselves.

In assessing the future impact of a technology, one must determine the likely policies of organizations and individuals. These policies include, first, whether to use a particular technology and, second, who will use it and how. I have argued that these policies are a function of not only the technology itself but also of the broader economic realities and, further, what I've termed as values. The overall picture is complex and merits further exploration.

Figure 3 is a diagram of the four factors, with arrows between them representing all possible directions of interaction. From this perspective, the main purpose of this article is to explore the arrow going from technology to policies, that is, to determine the effects of expert system technology on particular corporate or governmental policies (that is, job descriptions, qualifications, work processes, lines of communication, and so on). However, is there, in fact, a direct link between technology and policies, or is the effect always mediated by the other two factors, market and values? The following examples of the links that exist between the four factors in figure 3 serve to clarify this question and the matrix in general. (These examples are meant to illustrate each link and are not exhaustive).

Market $\Rightarrow$ Technology

In Advanced Manufacturing Technology, I established that new technologies are widely successful only when a market exists and the economics involved allow the development of such technology. Further, it's clear that the direction of development of existing tech-

nologies will likely be pulled by demand; an obvious example is the demand for cheaper, faster, easier-to-use microcomputers and the resultant more powerful central processing units, larger and cheaper memory devices, and the explosion of user-friendly software.

Technology $\Rightarrow$ Market

New technologies can create or shape new markets. As examples, the internal combustion engine and the subsequent popularization of the automobile gave birth to the tire industry and dozens of related businesses. The electric light bulb and audio recording made the film industry conceivable, and more recently, the microcomputer has provided the opportunity for software developers to flourish. New technologies can also destroy the demand for existing products: Witness the buggy whip and mechanical calculator.

Market $\Rightarrow$ Values

Can market forces affect values? Under this category, I suggest that the market, operating in the guise of advertising, certainly shapes our notions of such things as beauty, appropriate behavior, and worthy goals.

Values $\Rightarrow$ Market

This case refers to how our culture-specific values and ethics direct the price (or existence) of goods in the market. Simple supply and demand arguments only superficially explain the huge salaries paid to football players, \$500 per night hotel rooms, \$75 haircuts, the existence of psychotherapists for dogs, or the reason rabbits are raised and sold for human consumption in France but not in the United States.

Market $\Rightarrow$ Policies

A clear example on a microeconomic level is the situation of high demand for a product and suppliers or would-be suppliers attempting to satisfy this demand. A manufacturer might respond to this demand by enacting policies to increase production: hiring additional workers or buying new machines, for example. On a macroeconomic level, free market excesses might result in government policies (minimum wage laws, antitrust regulations, and so on) to curb these excesses.

Policies $\Rightarrow$ Market

A case in point is joint action by suppliers who form an oligopoly (for example, oil pro-

ducers). Possible policies include enforcing minimum price standards or restricting production. Government policies such as setting interest rates, commodity price supports, or tariffs on imported goods clearly have a serious impact on the market.

Technology $\Rightarrow$ Values

How does technology in and of itself shape the way we view the world? Two examples should make this relationship evident. The first is that modern technology tends to foster the belief that we are masters of nature and not part of it. Natural disasters are shocking not only because of the immediate toll on human life and property but also because they force us to realize that technology does not offer absolute protection from the elements. Second, the argument can be made that Western Europe's technical prowess in transportation and weaponry in the seventeenth, eighteenth, and nineteenth centuries not only made colonial expansion possible but also fostered the moral arguments (that is, the white man's burden) that sincere or not were supplied as ethical justification.

Values $\Rightarrow$ Technology

Here, the issue is how our value systems dictate the development and use of technology. The most direct examples in this area are personal decisions by scientists and engineers concerning the areas of research they will or will not engage in based on their moral convictions. This decision might relate to weapons or nuclear power research or something less dramatic; for example, Do I design computer systems to replace humans in their tasks or to assist humans in their tasks?

Values $\Rightarrow$ Policies

Laws and government regulations are the clearest cases of policies reflecting society's values. Widely accepted ethical rules of behavior also illustrate the point.

Policies $\Rightarrow$ Values

I return here to Taylorism and the policies that dictated assembly line behavior. The results of these policies, including bored, disinterested, and alienated workers, fed the belief that workers (or at least certain workers) are inherently bored, disinterested, and alienated. (Here is a case of a policy supporting values that helped create the policy in the first place.)

. . . modern technology tends to foster the belief that we are masters of nature and not part of it.

Policies $\Rightarrow$ Technology

In deciding where to spend their research and development funds, companies, research laboratories, and governments exemplify how policy decisions directly affect the direction of technological developments. These interactions are more complicated than the simple factor-to-factor effects described. Feedback can occur, as in the case of values affecting policies affecting values (described earlier) or market affecting policy affecting market (market excesses resulting in government policies causing other market problems that are dealt with by government policies, and so on). Further, it's likely that most interactions involve more than two factors; one can imagine values resulting in policies that have an impact on the market and, thus, direct technology.

However, I can think of no examples of technology directly dictating policy decisions. The mediating variables of values and market always have substantial impact. Technology provides options. In the manufacturing case, JIT, participatory management, increased flexibility, and so on, are simply not possible without the support of computer-based technologies. They are not, however, the inescapable, immediate result of the technology itself.

Two Scenarios

Given the matrix and its foundations, what can be said about the impact of expert system technology on work and the organization? I propose two scenarios.

Scenario 1: No Widespread Impact.

First, note that nowhere in the study of leading-edge manufacturers cited in the previous section is there a single word regarding expert systems. Underlying this scenario is the presumption that expert systems are simply not broadly applied because of their fundamental limitations, and until some breakthrough occurs, their drawbacks will continue to limit their use. Such breakthroughs might, for example, be along the fronts of common-sense reasoning, improved knowledge engineering, or reasoning from basic principles. Again, my assumption is that no such breakthrough will occur and be ready for commercialization in the three- to five-year time horizon of this exercise. With reference to the technology matrix, this scenario stipulates that a technology, without the pull of the market, has minimal effect on policy. This skepticism concerning the actual use of

expert systems is supported by a generally optimistic article in *Fortune* magazine (Kupfer 1987). It included the following comments:

Though expert systems are now in the mainstream computer world, they are not yet fully established. While American Express is proud enough of its expert system to have featured it in its last annual report, company officials are discomfited by the difficulty they are having in assimilating the new technology. People still regard expert systems as clever, industrious immigrants who wear funny clothes and can't quite be trusted.

I now relax the assumption of nonpenetration. Expert systems do, in fact, perform remarkably well on particular types of problems. The use of expert systems on these applications can indeed save large amounts of money and provide the benefits outlined earlier in this article. Examples of major companies that have developed expert systems for repair of engines and generators, credit evaluation, and the manufacture of automobile and airplane parts are impressive. One source (Feigenbaum, McCorduck, and Nii 1988) provides details on some 20 successful expert system applications and alludes to hundreds more. These applications are, typically, diagnosis-type tasks, although other problem types (monitoring or scheduling for example) are conceivable. Commonly, the expert system is designed to replicate the performance of highly skilled, experienced individuals, thereby allowing the task to be performed by less skilled, less experienced individuals in conjunction with the system. The experts are then free for other (hopefully more interesting) tasks. (The task might, in fact, involve enhancing the knowledge of the system.) With this in mind, I propose scenario 2.

Scenario 2: Expert Systems Will Contribute to and Accelerate the Impact Already Observed with Advanced Manufacturing Technologies.

Although expert systems are structurally—and, in many ways, functionally—different than AMTs, I argue (by analogy) that expert systems fit, in several fundamental ways with respect to their impact, into the type of dynamic environment described at the head of this section and introduced here under the rubric of AMTs.

The first factor is skill upgrading. Workers with lower skills take on jobs previously done by higher-skilled individuals. As noted previously, this decision is one of policy. Workers with varying skill levels might or might not

be afforded the time, incentives, and resources to think about process improvement. (See later comment regarding intellectual content.)

Second is flexibility. Training time is shorter; novices can quickly gain skill by working with the system. Further, “old hands” who have worked on another task for some time (remember the broad job specifications with AMTs) can fairly easily return to their previous work with the help of the system.

A third factor is the possibility of intellectual content. Expert systems are typically living and breathing systems; that is, they need to be frequently updated. Users can be part of the process. Clearly, they are likely to be in a good position to find deficiencies and suggest improvements in the system because they work with it. However, it also requires that they develop a sound understanding of the process in question. This understanding must, therefore, be provided through thoughtful design of the expert system. Zuboff (1985, 1988) refers to this type of system design as *informating*, as opposed to automating, and makes the argument that informating a process, as opposed to automating it, is a policy decision. Some expert systems for maintenance of large systems (for example, power plants) and those based on first principles tend to implicitly support user comprehension.

Fourth is interdependence. In a large organization, the development, diffusion, and updating process requires networking and effective communication across scattered (both functionally and geographically) areas of the organization. Expert systems can enhance knowledge and information management, a crucial function in the turbulent environment of the postindustrial organization.

Fifth is increased capital cost per worker. These observations are supported by current research (Weitz and DeMeyer 1989) into the development and diffusion of a large-scale expert system for computer board diagnosis. The system is designed for use by a multinational company across multiple sites throughout Europe.

Note that both these scenarios rest heavily on an attempt to understand the technology itself, the market for the technology, some general sense of values, and the policy decisions likely to result.

Observations and Questions

“Cheshire Puss,” she [Alice] began, “Would you tell me, please, which way I ought to go

Future research in the area of technology impact should be directed toward measuring and quantifying the factors in the technology matrix.

from here?”

“That depends a good deal on where you want to get to,” said the Cat.

—Lewis Carroll (as cited in Armstrong [1985])

In this article, I provided two scenarios for the likely impact of expert systems on the nature of work and the structure of the organization over the medium term. Moreover, a general framework was developed for the process of assessing the future impact of a technology. This framework can be summarized by the following four guidelines:

First, any forecast of the diffusion or impact of a technology should stress the market for the technology and not the technology itself.

Second, forecasts made by those in the business—in the case of expert systems, software and hardware developers and salespeople—should be viewed with deep suspicion. They are likely to be optimistic.

Third, all forecasts should explicitly state the assumptions on which they rest and, better yet, should evaluate the stability of these assumptions and the sensitivity of the forecasts to possible variations from these assumptions.

Fourth, the impact of a technology is a function of policies regarding how the technology is to be used. These policy choices are a function of the business environment; the culture in which the technology is introduced; and, indeed, the technology itself.

The first of these guidelines suggests that appropriate further work in impact assessment for expert systems should include fundamental market research. Although the characteristics of a problem that make it amenable to expert system solution have been isolated (Dym 1987; Bobrow, Mittal, and Stefik 1986; Weitz and DeMeyer 1989), the purpose of this market research should be to realistically determine the number of organizations with these problems and whether the costs, benefits, and potential strategic advantage afforded by an expert system solution

support or discourage the likelihood that expert system technology will be applied.

Future research in the area of technology impact should be directed toward measuring and quantifying the factors in the technology matrix. This approach would provide a sounder understanding of the factors in the matrix and how they interact as well as move the matrix in the direction of a substantive model with predictive validity. With such a model, the following types of questions could be addressed. First, which options resulting from new technologies are inescapable (if any), and when do they become so? Economic, political, and social pressure might force certain uses of new technologies to be, to one degree or another, more desirable than others. How are these pressures created and by whom? Second, when the market and values conflict, for example, when participatory management is not more effective than a Taylorized system, under what circumstances does one mediating factor take precedence over the other in directing the policies governing the acceptance and use of a technology? Further work exploring the nature of the interactions between technology, market, values, and policies is needed, not only to better forecast the impact of technology but also to help more coherently direct this impact.

Acknowledgments

I would like to thank Dr. Thomas Bernold of the Gottlieb Duttweiler Institute for suggesting the topic of this work. I would also like to thank Spyros Makridakis and Arnoud DeMeyer of INSEAD and Ted Stohr of New York University for their helpful comments on earlier drafts of this article.

References

- Armstrong, J. S. 1985. *Long Range Forecasting*. New York: Wiley.
- Bobrow, D. G.; Mittal, S.; and Stefik, M. J. 1986. Expert Systems: Perils and Promise. *Communications of the ACM* 29(9): 880–894.
- Dym, C. L. 1987. Issues in the Design and Implementation of Expert Systems. *Artificial Intelligence for Engineering Design, Analysis, and Manufacturing* (AI EDAM) 1(1): 37–46.
- Feigenbaum, E.; McCorduck, P.; and Nii, H. P. 1988. *The Rise of the Expert Company*. New York: Times.
- Harmon, P., and King, D. 1985. *Artificial Intelligence in Business: Expert Systems*. New York: Wiley.
- Hayes, R., and Wheelwright, S. 1984. *Restoring Our Competitive Edge: Competing through Manufacturing*. New York: Wiley.

Hayes-Roth, F.; Waterman, D.; and Lenat, D., eds. 1983. *Building Expert Systems*. Reading, Mass.: Addison-Wesley.

Hogarth, R., and Makridakis, S. 1981. Forecasting and Planning: An Evaluation. *Management Science* 27(2): 115–138.

Huber, G. 1984. The Nature and Design of Post-Industrial Organizations. *Management Science* 30(8): 928–951.

Kupfer, A. 1987. Now, Live Experts on a Floppy Disk. *Fortune* 116(8): 47–52.

Makridakis, S. 1988. Business Firms and Managers in the 21st Century. *Long Range Planning*. Forthcoming.

Makridakis, S. 1986. The Art and Science of Forecasting. *International Journal of Forecasting* 2(1): 15–39.

Makridakis, S.; Wheelwright, S.; and McGee, V. 1983. *Forecasting Methods and Applications*. New York: Wiley.

Markus, M. L., and Robey, D. 1988. Information Technology and Organizational Change: Causal Structure in Theory and Research. *Management Science* 34(5): 583–598.

Schnaars, S., and Berenson, C. 1986. Growth Market Forecasting Revisited: A Look Back at a Look Forward. *California Management Review* 27(4): 71–88.

Walton, R., and Susman, G. 1987. People Policy for the New Machines. *Harvard Business Review* 65(2): 98–107.

Weitz, R. 1986. NOSTRADAMUS: A Knowledge-Based Forecasting Advisor. *International Journal of Forecasting* 2(3): 273–283.

Weitz, R., and DeMeyer, A. 1989. Managing Expert Systems: A Framework and Case Study. Working paper, 89/45, INSEAD. Also in *Information and Management*. Forthcoming.

Zuboff, S. 1988. *In The Age of the Smart Machine: The Future of Work and Power*. New York: Basic.

Zuboff, S. 1985. Automate/Informate: The Two Faces of Intelligent Technology. *Organizational Dynamics* 14(2): 5–18.

Rob Weitz is a visiting associate professor of Information Systems at the Stern Graduate School of Business at New York University. He has served on the faculty of INSEAD, the European Institute of Business Administration, in Fontainebleau, France. Weitz's research focuses on expert systems, forecasting, and the interaction of technology and the organization.