UNCERTAINTY AND TECHNOLOGICAL CHANGE

Nathan Rosenberg*

I would like to begin with two generally accepted propositions: First, technological change is a major ingredient of long-term economic growth, and second, technological change is characterized by a high degree of uncertainty. Understanding the nature of these uncertainties and the obstacles to surmounting them is not a trivial matter. Rather, it goes to the heart of how new technologies are devised, how rapidly they diffuse, the ultimate extent of that diffusion, and their eventual impact on economic performance and welfare.

In view of the great uncertainties attached to the innovation process, it is hardly surprising that innovating firms have, historically, experienced high failure rates. Quite simply, the vast majority of attempts at innovation fail. But to describe the high failure rate associated with past innovation is to tell only a part of the story, and perhaps not the most interesting part. Indeed, I want to suggest that the more intriguing part of the story, with which I will be mainly concerned, has been the inability to anticipate the future impact of successful innovations, even after their technical feasibility has been established. This statement remains valid whether we focus on the steam engine 200 years ago or on the laser within our own lifetimes.

I will suggest that uncertainty is the product of several sources and that it has a number of peculiar characteristics that shape the innovation process and, therefore, the manner in which technological change exer-

---

*Fairleigh S. Dickinson, Jr., Professor of Public Policy, Department of Economics, Stanford University. The author wishes to acknowledge valuable comments on earlier drafts by Moses Abramovitz, Victor Fuchs, Ralph Landau, Roberto Mazzoleni, Richard Nelson, Richard Rosenbloom, Scott Stern, and members of the Program on Economic Growth and Policy of the Canadian Institute for Advanced Research.

cises its effects on the economy. Since I will be concerned primarily with what has shaped the trajectory and the economic impact of new technologies, my focus will be confined to technologies that have had significant economic consequences.

I should also say at the outset that, while I am not primarily concerned with the recent formal literature on growth theory (specifically The New Growth Theory), I am surprised that that literature has, so far at least, omitted any mention of uncertainty. While the rate of innovation is surely a function of the degree to which investors can appropriate the gains from their innovation, a number of central features of the innovation process revolve around uncertainty. At the very least, when evaluating projects, a risk/return trade-off that reflects the uncertainty attaching to appropriability must be considered. But the kinds of uncertainties that will be identified here go far beyond the issue of appropriability.

One further caveat seems appropriate. The discussion that follows is "anecdotal" in nature. However, the anecdotes have been deliberately selected to include many of the most important innovations of the twentieth century. Thus, if the characterizations offered below stand the test of further scrutiny, the analysis of this paper will have captured distinct features of the innovation process for technologies whose cumulative economic importance has been immense.

It is easy to assume that uncertainties are drastically reduced after the first commercial introduction of a new technology, and Schumpeter offered strong encouragement for making that assumption. His views have proven to be highly influential. In Schumpeter's world, entrepreneurs are compelled to make decisions under circumstances of very limited and poor quality of information. But in that world, the successful completion of an innovation resolves all the ex ante uncertainties. Once invention occurs, the stage is set for imitators, whose actions are responsible for the diffusion of a technology. Perhaps it should be said that the stage is now set for "mere imitators," for Schumpeter was fond of preceding the noun "imitators" with the adjective "mere." The point is one of real substance, and not just linguistic usage. In Schumpeter's view, life is easy for the imitators, because all they need to do is to follow in the footsteps of the entrepreneurs who have led the way, and whose earlier activities have resolved all the big uncertainties.

It is, of course, true that some uncertainties have been reduced at that point. However, after a new technological capability has been established, the questions change and, as we will see, new uncertainties, particularly uncertainties of a specifically economic nature, begin to assert themselves.

The purpose of this paper is to identify and to delineate a number of important aspects of uncertainty as they relate to technological change. These aspects go far beyond those connected with the inventive process alone. In addition, they reflect a set of interrelated forces that are at the heart of the relationship between changes in technology and improvements in economic performance.

Some Historical Perspectives

Consider the laser, an innovation that is certainly one of the most powerful and versatile advances in technology in the twentieth century, and one that is surely still in the early stages of its trajectory of development. Its range of uses in the 30 years since it was invented is truly breathtaking, and would include precision measurement, navigation, and chemical research. It is also essential for the high-quality reproduction of music in compact discs (CDs). It has become the instrument of choice in a range of surgical procedures, including extraordinarily delicate eye surgery, where it is used to repair detached retinas, and gynecological surgery, where it now provides a simpler and less painful method for removal of certain tumors. It is extensively employed in gallbladder surgery. The pages of this manuscript were originally printed by a laser (Hewlett Packard laser jet printer). It is widely used throughout industry, including textiles, where it is employed to cut cloth to desired shapes, and metallurgy and composite materials, where it performs similar functions. The opening sentence in an article appearing in *The New York Times* in April 1996 stated: "Lawrence Livermore Laboratory, builder of lasers powerful enough to shoot down missiles or ignite miniature hydrogen bombs, has created a portable laser that is said to be able to obliterate graffiti from walls and statues at lightning speed."

But perhaps no single application of the laser has been more profound than its impact on telecommunications where, together with fiber optics, it is revolutionizing transmission. The best transatlantic telephone cable in 1966 could carry simultaneously only 138 conversations between Europe and North America. The first fiber optic cable, installed in 1988, could carry 40,000. The fiber optic cables being installed in the early 1990s can carry nearly 1.5 million conversations (Wriston 1992, pp. 43–44). And yet it is reported that the patent lawyers at Bell Labs were initially unwilling even to apply for a patent on the laser, on the grounds that such an invention had no possible relevance to the telephone industry. In the words of Charles Townes, who subsequently won a Nobel Prize for his research on the laser, "Bell's patent department at first refused to patent our amplifier or oscillator for optical frequencies because, it was explained, optical waves had never been of any importance to communications and hence the invention had little bearing on Bell System interests" (Townes 1968, p. 701).

Let me cite some further major historical instances where the common theme is the remarkable inability, at least from a later perspective, to foresee the uses to which new technologies would soon be put. Western Union, the telegraph company, was offered the opportunity to purchase Bell's 1876 telephone patent for a mere $100,000, but turned it down. In fact, "Western Union was willing to withdraw from the telephone field in 1879 in exchange for Bell's promise to keep out of the telegraph business." But if the proprietors of the old communications technology were myopic, so too was the patent holder of the new technology. Alexander Graham Bell's 1876 patent did not mention a new technology at all. Rather, it bore the glaringly misleading title "Improvements in Telegraphy" (Brock 1982, p. 90).

Marconi, who invented the radio, anticipated that it would be used primarily to communicate between two points where communication by wire was impossible—as in ship-to-ship or ship-to-shore communication. (To this day the British call the instrument the "wireless," precisely reflecting Marconi's early conceptualization.) Moreover, the radio in its early days was thought to be of potential use only for private communication: that is, point-to-point communication, rather like the telephone, and not at all for communicating to a large audience of listeners. Surprising as it may seem to us today, the inventor of the radio did not think of it as an instrument for broadcasting. Marconi, in fact, had a conception of the market for radio that was the precise opposite of the one that actually developed. He visualized the users of his invention as steamship companies, newspapers, and navies that required directional, point-to-point communication—"narrowcasting" rather than broadcasting. The radio should therefore be capable of transmitting over great distances, but the messages should be private, not public (Douglas 1987, p. 34).

The failure of societal imagination was widespread. According to one authority: "When broadcasting was first proposed ... a man who was later to become one of the most distinguished leaders of the industry announced that it was very difficult to see uses for public broadcasting. About the only regular use he could think of was the broadcasting of Sunday sermons, because that is the only occasion when one man regularly addresses a mass public" (Martin 1977, p. 11).

The wireless telephone, when it became feasible in the second decade of the twentieth century, was thought of in precisely the same terms as the wireless radio. J.J. Carty, who was chief engineer of the New York Telephone Company, stated in 1915, "The results of long-distance tests show clearly that the function of the wireless telephone is primarily to reach inaccessible places where wires cannot be strung. It will act mainly as an extension of the wire system and a feeder to it" (Maclaurin 1949, pp. 92–93).

The computer, in 1949, was thought to be of potential use only for rapid calculation in a few scientific research or data processing contexts. The notion of a large potential market was rejected by no less a person than Thomas Watson, Sr., at the time the president of IBM. The prevailing view before 1950 was that world demand probably could be satisfied by just a few computers (Ceruzzi 1987, pp. 188–93).

The invention of the transistor, certainly one of the greatest inventions of the twentieth century, was not announced on the front page of *The New York Times*, as might have been expected, when it was made public in December 1947. On the contrary, it was a small item buried deep in the newspaper's inside pages, in a regular weekly column titled "News of Radio." It was suggested there that the device might be used to develop better hearing aids for the deaf, but nothing more.

This listing of failures to anticipate future uses and larger markets for new technologies could be expanded almost without limit. We could, if we liked, amuse ourselves indefinitely at the failure of earlier generations to see the obvious, as we see it today. But that would be a mistaken conceit. For reasons that I propose to examine, I am not particularly optimistic that our ability to overcome the ex ante uncertainties connected with the uses of new technologies is likely to improve drastically. If I am right, a more useful issue to explore is what incentives, institutions, and policies are more likely to lead to a swifter resolution of these uncertainties.

Much of the difficulty, I suggest, is connected to the fact that new technologies typically come into the world in a very primitive condition. Their eventual uses turn upon an extended improvement process that vastly expands their practical applications. Thomas Watson, Sr., was not necessarily far off the mark when he concluded that the future market for the computer was extremely limited, *if one thinks of the computer in the form in which it existed immediately after the Second World War*. The first electronic digital computer, the ENIAC, contained no less than 18,000 vacuum tubes and filled a huge room. (It was more than 100 feet long.) Any device that has to rely on the simultaneous working of 18,000 vacuum tubes is bound to be notoriously unreliable. The failure in prediction was a failure to anticipate the demand for computers after they had been made very much smaller, cheaper, and more reliable, and when their performance characteristics, especially their calculating speed, had been improved by many orders of magnitude. That is to say, the failure was the inability to anticipate the trajectory of future improvements and the economic consequences of those improvements.

If space permitted, the history of commercial aviation could be told in similar terms, as could the history of many other innovations. With respect to the introduction of the jet engine, in particular, the failure to anticipate the importance of future improvements occurred even at the most eminent scientific levels. In 1940, a committee of the National Academy of Sciences was formed to evaluate the prospects for developing a gas turbine for aircraft. The committee concluded that such a turbine was quite impractical because it would have to weigh 15 pounds for each horsepower delivered, whereas existing internal combustion engines weighed only slightly over one pound for each horsepower delivered. In fact, within a year the British were operating a gas turbine that weighed a mere four-tenths of one pound per horsepower (U.S. Navy, Bureau of Ships 1941, p. 10).

This is an appropriate place at which to make a very simple, but nonetheless fundamental observation: Most R&D expenditures are devoted to product improvement. According to McGraw-Hill annual surveys over a number of years, the great bulk of R&D (around 80 percent) is devoted to improving products that already exist, rather than to the invention of new products. Thus, it is incorrect to think of R&D expenditures as committed to the search for breakthrough innovations of the Schumpeterian type. On the contrary, the great bulk of these expenditures need to be thought of as exhibiting strongly path-dependent characteristics. Their main goal is to improve upon the performance of technologies that have been inherited from the past.

A moment's reflection suggests that this should not be surprising. The telephone has been around for more than a hundred years, but only recently has its performance been significantly enhanced by facsimile transmission, electronic mail (e-mail), voice mail, data transfer, on-line services, mobile phones, conference calls, and "800" numbers. The automobile and the airplane are each more than 90 years old, the camera is 150 years old, and the Fourdrinier machine, which is the mainstay of the papermaking industry today, was patented during the Napoleonic Wars. Clearly the improvement process deserves far more attention than is suggested by Schumpeter's frequent recourse to the derisory term "mere imitators." Equally clearly, a world in which most R&D expenditures are devoted to improving upon technologies that already exist is also a world in which technological change can hardly be characterized as exogenous.

So far it has been suggested, by citing important historical cases, that uncertainty plays a role in technological change that goes far beyond the uncertainty associated with technological feasibility alone. Indeed, the uncertainty associated with the eventual uses of the laser or the computer might, more appropriately, be characterized as "ignorance" rather than as "uncertainty." That is to say, along any particular dimension of uncertainty, decisionmakers do not have access to an even marginally informative probability distribution with respect to potential outcomes. It is not difficult to demonstrate that ignorance plays a large part in the process of technological change! However, rather than arguing over the differences between Arrovian and Knightian uncertainty, the next section of this paper will outline a number of important dimensions along which uncertainty plays a role in the rate and direction of inventive activity and diffusion. Taken together, we have very little information, *even retrospectively,* about the relationships among these different dimensions. If uncertainty exists along more than one dimension, and the decisionmaker does not have information about the joint distribution of all the relevant random variables, then we have little reason to believe that a "rational" decision is possible or that a well-defined "optimal" investment or adoption strategy will be found.

The Dimensions of Uncertainty

Why is it so difficult to foresee the impact of even technologically practicable inventions? Much of the relevant literature emphasizes the huge uncertainty associated with the question: "Will it work?" This is clearly a major source of uncertainty, but the fixation upon workability has served to distract attention from several other, more subtle and overlapping sources. We turn now to a consideration of these sources.

Ex Ante Uncertainty about Improvements and Uses

It is not only that new technologies come into the world in a very primitive condition; they often do so with properties and characteristics whose usefulness cannot be immediately appreciated. It is inherently difficult to identify uses for new technologies. The laser (Light Amplification by Stimulated Emission of Radiation) represents, at one level, simply a light beam formed by the excitation of atoms at high energy levels. It has turned out that laser action can occur with a wide range of materials, including gases, liquids, and solids. The uses to which this capability has been put have been growing for 30 years, as suggested earlier, and will doubtless continue to grow for a long time, just as it took many decades to explore the uses to which electricity could be put after Faraday discovered the principles of electromagnetic induction in 1831.[1]

An essential aspect was that neither the laser nor electricity represented an obvious substitute for anything that already existed. Neither had a clearly defined antecedent. Rather, each technology was a newly discovered phenomenon that was the outcome of pure scientific research.[2]

In the field of medical diagnostics it has frequently happened that, after some new visualization technology has been developed, it has taken

[1] It is recorded that a skeptical MP turned up at Faraday's laboratory shortly after his discovery of electromagnetic induction and asked him in a rather supercilious tone what it was good for. Faraday is supposed to have replied: "Sir, I do not know what it is good for. But of one thing I am quite certain: someday you will tax it."

[2] In fact, Einstein had already worked out the pure science underlying laser action in 1916, in a paper on stimulated emission. From the point of view of the history of science, it might be said that there was "nothing new" when laser technology was developed some 45 years later, although a Nobel Prize was awarded for the achievement. From the point of view of technological change and its economic and social impact, the development of the laser was of course a major event.

a long time to learn how to translate the new observational capability into clinically useful terms. This has been the case with respect to CAT scanners, magnetic resonance imaging (MRI), and most recently echocardiography. Often a great deal of time-consuming additional research has been required before it was possible to make a reliable, clinically helpful interpretation of what was already being visualized in terms of the diagnosis of a disease condition in the heart, lungs, or brain.

This is presently the case with respect to PET—positron emission tomography. PET scanners are powerful tools for providing a quantitative analysis of certain physiological functions, unlike CAT and MRI, which are valuable for anatomical observation. Thus, it has great potential for providing useful information on the effectiveness, for example, of drug therapy for the treatment of various diseases, such as brain tumors. But, quite aside from the huge cost of this technology, its clinical application in such fields as neurology, cardiology, and oncology has so far been limited by the continuing difficulties of translating observations and measurements of physiological functions into specific, meaningful clinical interpretations.

A related point can be made in the currently burgeoning field of medical innovation. The inherent complexity of the human body and, perhaps equally important, the heterogeneity of human bodies, have rendered it extremely difficult to tease out cause–effect relationships, even in the case of medications that have been widely used for long periods of time. Aspirin (acetylsalicylic acid), probably the world's most widely used drug, has been in use for very nearly a century, but only in the last couple of years has its efficacy been established for reducing the incidence of heart attacks as a consequence of its blood-thinning properties.

Although the discovery of negative side effects has received far more public attention, the discovery of unexpected beneficial new uses for old pharmaceutical products is a common, and often serendipitous, experience. Another significant case in point has been the applications of andrenergic beta-blocking drugs, one of the more significant medical innovations of our time. These compounds were originally introduced for the treatment of two cardiovascular indications, arrythmias and angina pectoris. Today they are used in the treatment of more than 20 diverse conditions, largely as a result of new uses that were uncovered after they had been introduced into cardiology. These include such noncardiac indications as gastrointestinal bleeding, hypertension, and alcoholism (Gelijns 1991, pp. 121 and 269). Similar experiences could be related with respect to AZT (currently employed in the treatment of AIDS patients), oral contraceptives, RU-486, streptokinase, alpha interferon, and Prozac. More generally, the widespread "off-label" uses of many drugs provide a good indication of the pervasiveness of ex ante uncertainty in medical innovation.

The Need for Complementary Technologies

Second, the impact of an innovation depends not only on improvements of the invention, but also upon improvements that take place in complementary inventions. For the lawyers at Bell Labs to have had some appreciation of the laser's importance for telephone communication, they would have required some sense of fiber optic technology and the ways in which the two—lasers and fiber optics—might be combined. The laser was of no particular use in telephone transmission without the availability of fiber optics. Telephone transmission is being transformed today by the *combined* potential of these two technologies. Optical fiber did in fact exist in its own rather primitive form in the early 1960s, when the first lasers were developed, but not in a form that could accommodate the requirements of telephone transmission. In fact, it is interesting to note that an excellent book on the telecommunications industry, published as recently as 1981, provides no discussion whatsoever of this new fiber optic technology (Brock 1982). As is often the case, it took a number of years for some of the attractive properties of fiber optics technology to become apparent: the lack of electromagnetic interference, the conservation of heat and electricity, and the enormous expansion in bandwidth that fiber optics can provide—the last feature a consequence of the fact that the light spectrum is approximately 1,000 times wider than the radio spectrum.

The general point is that the impact of invention A will often depend upon invention B, and invention B may not yet exist. But perhaps a more useful formulation is to say that inventions will often give rise to a search for complementary inventions. An important impact of invention A is to increase the demand for invention B. The declining price of electricity, after the introduction of the dynamo in the early 1880s, stimulated the search for technologies that could exploit this unique form of energy. But the time frame over which such complementary innovations could be developed turned out to vary considerably. The search gave rise almost instantly to a burgeoning electrochemical industry, employing electrolytic techniques (aluminum), but a much longer period of time was required before the development of the complementary electric motor that was to become ubiquitous in the twentieth century.

Similarly, a main reason for the modest future prospects that were being predicted for the computer in the late 1940s was that transistors had not yet been incorporated into the computers of the day. The introductions of the transistor, and later integrated circuits, into computers were, of course, momentous events that transformed the computer industry. Indeed, in one of the most remarkable technological achievements of the twentieth century, the integrated circuit eventually *became* a computer with the advent of the microprocessor in 1970. The world would be a far different place today if computers were still made using vacuum tubes.

The need to develop complementary technologies may have had a great deal to do in the last couple of decades with the apparent failure of computer technology to raise the level of productivity growth in the United States above its recent rather dismal levels. Robert Solow has made the observation that we see computers everywhere today except in the productivity statistics. But it appears to be typical of truly major innovations that they take a long time to absorb. The historical experience with respect to the introduction of electricity offers many earlier parallels. If we date the beginning of the electric age in the early 1880s (dynamos), it was fully 40 years—into the 1920s—before the electrification of factories began to show up in terms of significant measured productivity growth (Du Boff 1967; Devine 1983; Schurr 1990).

Major new technological regimes take many years before they replace an established technology. Partly the delay is due to having to develop numerous components of a larger technological system, an issue that will be addressed shortly. Restructuring a factory around an electric power source, in place of the earlier steam engine or water power, commonly required a complete redesign and restructuring of a factory facility. It represented, among other things, a revolution in the principles of factory organization. The layout of the machinery in the factory now had far more flexibility than it did with the old power sources. Learning how best to exploit a new, highly versatile power source with entirely different methods of power transmission inside the plant involved decades of experimentation and learning. Indeed, such technological innovations commonly require significant organizational changes as well. (The glacial pace at which organizational changes often take place may have a great deal to do with Solow's complaint about the failure of computers to be reflected in the productivity statistics.)

Moreover, firms that had huge investments in manufacturing plants, with long productive lives still ahead of them, naturally were reluctant to discard a facility that was still perfectly usable. As a result, if we ask who the early adopters of electricity were in the first 20 years of the twentieth century, it turns out that they were mainly new industries that were setting up production facilities for the first time; that is, producers of "tobacco, fabricated metals, transportation equipment and electrical machinery itself." In the older, established industries, the introduction of electric power had to await the "physical depreciation of durable factory structures," and the "obsolescence of older-vintage industrial plants sited in urban core areas" (David 1990, p. 357).

The general point is that a radical new technology such as a computer must necessarily have a very long gestation period before its characteristics and opportunities are well understood and can be thoroughly exploited. In 1910 only 25 percent of U.S. factories used electric power. But 20 years later the figure had risen to 75 percent. History suggests that we should not be terribly surprised. Yet if we date the beginning of the modern computer—a much more complex general purpose technology than electricity—from the invention of the microprocessor in 1970, we are still only a quarter century into the computer age. It took some 40 years or so before electric power came to play a dominating role in manufacturing. History strongly suggests that technological revolutions are never completed overnight. If this is correct, it should be a source of optimism. The great economic benefits of the computer may still lie before us!

Innovations as Components of a Technological System

As a closely connected point, major technological innovations often constitute entirely new technological systems. But it is difficult in the extreme to conceptualize an entirely new system. Thus, thinking about new technologies is likely to be severely handicapped by the tendency to view them in terms of the old technologies that they eventually replace. Time and again, contemporaries of a new technology are found to have thought about it as a mere supplement that would offset certain inherent limitations of an existing technology. In the 1830s and 1840s, railroads were thought of merely as feeders into the existing canal system, to be constructed in places where the terrain had rendered canals inherently impractical (Fogel 1964). This is precisely the same difficulty that later was encountered by the radio. Similarly, the telephone was originally conceptualized as primarily a business instrument, like the telegraph, to be used to exchange very specific messages, such as the terms of a prospective contractual agreement. This may of course explain why Bell's telephone patent was, as mentioned earlier, titled "Improvements in Telegraphy."

It is characteristic of a system that performance improvements in one part are of only limited significance without simultaneous improvements in other parts. In this sense, technological systems may be thought of as comprising clusters of complementary inventions. Improvements in power generation can have only a limited impact on the delivered cost of electricity until improvements are made in the transmission network and the cost of transporting electricity over long distances. This need for further innovation in complementary activities is an important reason why even apparently spectacular breakthroughs usually have only a slowly rising productivity curve flowing from them. Within technological systems, therefore, major improvements in productivity seldom flow from single technological innovations, however significant they may appear to be. At the same time, the *cumulative* effects of large numbers of improvements within a technological system eventually may be immense.

Unanticipated Applications

An additional and historically very important reason why it has been so difficult to foresee the uses of a new technology is that many major inventions had their origins in the attempt to solve very specific and often very narrowly defined problems. However, it is common that once a solution has been found, it turns out to have significant applications in totally unanticipated contexts. That is to say, much of the impact of new technologies is realized through intersectoral flows. Inventions have very serendipitous life histories (Rosenberg 1976a).

The steam engine, for example, was invented in the eighteenth century specifically as a device for pumping water out of flooded mines. In fact it was, for a long time, regarded exclusively as a pump. A succession of improvements later rendered it a feasible source of power for textile factories, iron mills, and an expanding array of industrial establishments. In the course of the early nineteenth century, the steam engine became a generalizable source of power and had major applications in transportation—railroads, steamships, and steamboats. In fact, before the Civil War, the main use of the steam engine in the United States was not in manufacturing but in transportation. Later in the nineteenth century the steam engine was, for a time, used to produce a new and even more generalizable source of power—electricity—which in turn satisfied innumerable final uses to which steam power itself was not directly applicable. Finally, the steam turbine displaced the steam engine in the generation of electric power, and the special features of electricity—its ease of transmission over long distances, the capacity for making power available in "fractionalized" units, and the far greater flexibility of electricity-powered equipment—sounded the eventual death knell of the steam engine itself.

Major innovations such as the steam engine, once they have been established, have the effect of inducing further innovations and investments over a wide frontier. Indeed, the ability to induce such further innovations and investments is a reasonably good definition of what constitutes a major innovation. It is a useful way of distinguishing between technological advances that are merely invested with great novelty from advances that have the potential for a major economic impact. But this also highlights the difficulties in *foreseeing* the eventual impact, since that will depend on the size and the direction of these future complementary innovations and associated investments.

The life history of the steam engine was shaped by forces that could hardly have been foreseen by British inventors who were working on ways of removing water from increasingly flooded coal mines in the eighteenth century. Nevertheless, the very existence of the steam engine, once its operating principles had been thoroughly understood, served as a powerful stimulus to other inventions.

Impacts on Other Industries

I have been stressing here that innovations often arise as solutions to highly specific problems in a particular industry, and that their subsequent interindustry flow is bound to be highly uncertain. This is because the uses of a new technology in a quite different industrial context are especially difficult to anticipate. Moreover, in some cases a new technological capability may have multiple points of impact on another industry.

Consider the impact of the computer upon the air transportation industry. I would suggest that the changing performance of commercial air transportation has been at least as much influenced by the application of the computer to new uses in this industry as by the R&D spending that has taken place within air transportation itself.

- Supercomputers now perform a good deal of fundamental aerodynamic research, including much—but not all—of the research that was formerly performed in wind tunnels.
- Computers have been a major source of cost reduction in the design of specific components of the aircraft, such as the wing. They played an important role in the wing designs of the Boeing 747, 757, and 767, as well as the Airbus 310.
- Computers are now responsible for much of the activity that takes place in the cockpit, including of course the automatic pilot.
- Computers, together with weather satellites, which routinely determine the shifting location of high-altitude jet streams, are now widely used in determining optimal flight paths. The fuel savings for the world commercial airline industry is probably well in excess of $1 billion per year. (Note that this is yet another important case of the economic impact of a technology, the computer, depending upon a complementary technology that was only developed many years later, weather satellites.)
- Computers and computer networks are at the heart of the present worldwide ticketing and seating reservation system.
- Computer simulation is now the preferred method of instruction in teaching neophytes how to fly.
- The computer, together with radar, has become absolutely central to the operation of the air traffic control system, which would be difficult to conceive without it.

One important implication of this discussion is that R&D spending tends to be highly concentrated in a small number of industries. However, each of these few industries needs to be regarded as a locus of research activity that generates new technologies that may be widely diffused throughout the entire economy. Historically, a small number of industries have played this role in especially crucial ways: steam engines, electricity, machine tools, computers, transistors, and so on. This reinforces the earlier suggestion that we may even *define* a major—or breakthrough—innovation as one that establishes a new framework for the working out of incremental innovations. In this sense, incremental innovations are the natural complements of breakthrough innovations. Breakthrough innovations, in turn, have often provided the basis for the emergence of entirely new industries.

The Identification of Needs

The final constraint is rather less precise than the rest but, I believe, it is no less important. That is, the ultimate impact of some new technological capability is not just a matter of technical feasibility or improved technical performance; rather, it is a matter of identifying certain specific categories of human needs and catering to them in novel or cost-effective ways. New technologies need to pass an economic test, not just a technological one. Thus, the Concorde is a spectacular success in terms of flight performance, but it has proved to be a financial disaster, costing British and French taxpayers the equivalent of several billions of dollars.

Ultimately, what is often needed is not just technical expertise but the exercise of imagination. Understanding the technical basis for wireless communication, which Marconi did, was a very different matter from anticipating how the device might be used to enlarge the human experience. Marconi had no sense of this. On the other hand, an uneducated Russian immigrant, David Sarnoff, had a lively vision of how the new technology might be used to transmit news, music, and other forms of entertainment and information into every household (and eventually automobile) in the country. Sarnoff, in brief, appreciated the commercial possibilities of the new technology. Sarnoff's vision eventually prevailed, under his leadership of RCA after the First World War (Bilby 1985).

Similarly, Howard Aiken, a Harvard physics instructor who was a great pioneer in the early development of the computer, continued to think of it in the narrow context in which its early development took place—that is, purely as a device for solving esoteric scientific problems. As late as 1956 he stated: "if it should ever turn out that the basic logics of a machine designed for the numerical solution of differential equations coincide with the logics of a machine intended to make bills for a department store, I would regard this as the most amazing coincidence that I have ever encountered" (Ceruzzi 1987, p. 197). That is, of course, precisely how it turned out, but it was hardly a coincidence. A technology originally invented for one specific purpose—the numerical solution of large sets of differential equations—could readily be redesigned to solve problems in entirely different contexts, such as department store billing procedures. But it obviously was not obvious!

The essential point, of course, is that social change or economic impact is not something that can be extrapolated out of a piece of hardware. New technologies, rather, need to be conceived of as building blocks. Their eventual impact will depend on what is subsequently designed and constructed with them. New technologies are unrealized potentials that may take a very large number of eventual shapes. What shapes they actually take will depend on the ability to visualize how they might be employed *in new contexts*. Sony's development of the Walkman is a brilliant example of how an existing technological capability, involving batteries, magnetic tapes, and earphones, could be recombined to create an entirely new product that could provide entertainment in contexts where it previously could not be delivered—where, indeed, no one had even *thought* of delivering it—for example, to walkers or even joggers. To be sure, the product required a great deal of engineering redesign of existing components, but the real breakthrough was the identification, by Akio Morita, of a market opportunity that previously had not been identified.

Although many Americans continue to believe that the VCR was an American invention, that is simply an unsupportable perception. The American pioneers in this field, RCA and Ampex, gave up long before a truly usable product had been developed. Matsushita and Sony, on the other hand, made thousands of small improvements in design and manufacturing after the American firms had essentially left the field. These developments were closely connected to another point. A crucial step forward in the development of the VCR was the realization that a potential mass market existed in households if certain performance characteristics, especially the product's storage capacity, could be sufficiently expanded. Although the initial American conception of the VCR had been of a capital good to be used by television stations, some American as well as Japanese developers were aware of the much larger home market possibilities. The crucial difference seems to have been the Japanese confidence, based upon their own manufacturing experience, that they could achieve the necessary cost reductions and performance improvements. The rapid transformation of the VCR into one of Japan's largest export products was therefore an achievement of both imagination and justified confidence in their engineering capabilities (Rosenbloom and Cusumano 1987).

The limited view once held by Americans of the potential for the VCR bears some parallels with the disdain of the mainframe computer makers toward the personal computer as it began to emerge about 15 years ago. It was then fashionable to dismiss the PC as a mere "hacker's toy," with no real prospects in the business world and therefore no serious threat to the economic future of mainframes (*The New York Times* 1994).

Reviving Old Technologies, or Killing Them Off?

My analysis has focused upon barriers to the exploitation of new technologies. But in highly competitive societies with strong incentives to innovation, those incentives apply to improving old technologies as well as to inventing new ones. In fact, innovations often appear to induce vigorous and imaginative responses on the part of firms that find themselves confronted with close substitutes for their traditional products. It is not at all uncommon to find that the competitive pressure resulting from a new technology leads to an accelerated improvement in the old technology. Some of the greatest improvements in wooden sailing ships took place between 1850 and 1880, just after the introduction of the iron hull steamship and the compound steam engines that were to displace sailing ships by the beginning of the twentieth century. These innovations included drastic improvements in hull design that allowed greater speed, more cargo in proportion to the tonnage of the ship and, above all, the introduction of labor-saving machinery that reduced crew requirements by no less than two-thirds. Similarly, the greatest improvements in gas lamps used for interior lighting occurred shortly *after* the introduction of the incandescent electric light bulb (Rosenberg 1976b). More recently, soon after the introduction of coronary angioplasty, a potential substitute for coronary bypass surgery, substantial improvements were made in the "old" surgical procedure. In each case, of course, the timing may have been coincidental.

A major feature of the postwar telecommunications industry is that research has increased the capabilities of the already installed transmission system, in addition to leading to the development of new and more productive technologies. Every major transmission system—a pair of wires, coaxial cables, microwaves, satellites, fiber optics—has been subject to extensive later improvements in message-carrying capabilities, often with only relatively minor modification of the existing transmission technology. In some cases, order-of-magnitude increases have occurred in the message-carrying capability of an existing channel, such as a 3/8th inch coaxial cable, and such productivity improvements have frequently led to the postponement of the introduction of new generations of transmission technologies. For example, time-division multiplexing allowed an existing pair of wires to carry 24 voice channels or more, rather than the single channel that it originally carried. The same pattern is observed in fiber optics technology. When AT&T began field trials with fiber optics in the mid 1970s, information was transmitted at 45 megabytes per second. By the early 1990s, the standard for new fiber cables had reached 565 megabytes per second, with reliable sources predicting capacities of nearly 1,000 megabytes per second in the near future.

But it is not only the case that the introduction of new technologies often has to await the availability of complementary technologies and that, in the meantime, established technologies may achieve renewed competitive vigor through continual improvements. New technologies may also turn out to be substitutes rather than complements for existing ones, thus drastically shortening the life expectancy of technologies that once seemed to warrant distinctly bullish expectations. The future prospects for communication satellites declined quite unexpectedly during the 1980s with the introduction of fiber optics and the huge and reliable expansion of channel capacity that they brought with them. In turn, fiber optics, whose first significant application was in medical diagnostics in the early 1960s, may now be approaching the beginning of the end of its useful life. Fiber optic endoscopes had made possible a huge improvement in minimally invasive techniques for visualizing the gastrointestinal tract. Recently, new sensors from the realm of electronics, charged couple devices (CCDs), have begun to provide images with a resolution and degree of detail that could not possibly be provided by fiber optic devices. The CT scanner, certainly one of the great diagnostic breakthroughs of the twentieth century, is giving way to an even more powerful diagnostic capability—MRI. Uncertainties of this sort impart a large element of risk to long-term investments in expensive new technologies. The competitive process that eventually resolves these uncertainties is not the traditional textbook competition among producers of a homogeneous product, each seeking to deliver the same product to market at a lower cost. Rather, it is a competition among different technologies, a process that Schumpeter appropriately described as "creative destruction." Thus, it is no paradox to say that one of the greatest uncertainties confronting new technologies is the invention of yet newer ones.

The simultaneous advance in new technology, along with the substantial upgrading of old technology, underlines the pervasive uncertainty confronting industrial decisionmakers in a world of rapid technological change. One would have to be very optimistic, as well as naive, to think that some intellectual paradigm can be developed to handle all the relevant variables in some neat and systematic way. But it may be plausible to believe that a more rigorous analysis of the issues raised here may lead to considerable improvement in the way we think about the innovation process.

We can now return to the initial point: The lack of knowledge about the relationships between these different dimensions of uncertainty precludes us from understanding the total effect of uncertainty upon technological change. For example, two dimensions of uncertainty, discussed above, concern the refinement of complementary technologies and the potential for any technology to form the core of a new technological system. Even at the simplest level, it is difficult to be precise about the interaction between these different effects. The existence and refinement of complementary technologies may exercise a coercive and conservative effect, forcing the novel technology to be placed inside the current "system." Alternatively, however, complementary technologies may be exactly what is necessary for the practical realization of an entirely new system. My point is not to decide one way or the other on these issues; instead, it is to argue that a research program that neglects these interactions may be missing a very large part of how uncertainty has shaped the rate and direction of technological change and, by extension, the historical growth experience.

Concluding Observations

It is not part of my warrant to offer any policy recommendations. However, a few closing observations may be in order. The research community is currently being exhorted with increasing force to unfurl the flag of "relevance" to social and economic needs. The burden of much that has been said here is that frequently we simply do not *know* what new findings may turn out to be relevant, or to what particular *realm* of human activity that relevance may eventually apply. Indeed, I have been staking the broad claim that a pervasive uncertainty characterizes not only basic research, where it is generally acknowledged, but the realm of product design and new product development as well—the D of R&D. Consequently, early precommitment to any specific, large-scale technology project, as opposed to a more limited, sequential decision-making approach, is likely to be hazardous—that is, wasteful. Evidence for this assertion abounds in such government-sponsored projects as weapons procurement, the space program, research on the development of an artificial heart, and synthetic fuels.

The pervasiveness of uncertainty suggests that the government ordinarily should resist the temptation to play the role of a champion of any one technological alternative, such as nuclear power, or any narrowly concentrated focus of research support, such as the "War on Cancer." Rather, it would seem to make a great deal of sense to manage a deliberately diversified research portfolio, a portfolio that is likely to illuminate a range of alternatives in the event of a reordering of social or economic priorities or the unexpected failure of any single, major research thrust. Government policy ought to open many windows and provide the private sector with financial incentives to explore the technological landscape that can only be faintly discerned from those windows. Thus, my criticism of the federal government's postwar energy policy is not that it made a major commitment to nuclear power that subsequently turned out to be problem-ridden. A more appropriate criticism is aimed at the single-mindedness of the focus on nuclear power that led to a comparative neglect of many other alternatives, including not only alternative energy sources but improvements in the efficiency of energy utilization.

The situation with respect to the private sector is obviously different. Private firms may normally be expected to allocate their R&D funds in ways that they hope will turn out to be relevant. Private firms are very much aware that they confront huge uncertainties in the marketplace, and they are capable of making their own assessments and placing their "bets" accordingly. Bad bets are, of course, common, indeed so common that it is tempting to conclude that the manner in which competing firms pursue innovation is a very wasteful process. Such a characterization would be appropriate were it not for a single point: uncertainty. In fact, a considerable virtue of the marketplace is that, in the face of huge ex ante uncertainties concerning the uses of new technological capabilities, it encourages exploration along a wide variety of alternative paths. This is especially desirable in the early stages, when uncertainties are particularly high and when individuals with differences of opinion (often based upon differences in access to information) need to be encouraged to pursue their own hunches or intuitions. Indeed, it is important that this point be stated more affirmatively: The achievement of technological progress, in the face of numerous uncertainties, *requires* such ex ante differences of opinion.

Finally, a further considerable virtue of the marketplace is that it also provides strong incentives to terminate, quickly and unsentimentally, directions of research whose once rosy prospects have been unexpectedly dimmed by the availability of new data, by some change in the economic environment, or by a restructuring of social or political priorities. For a country that currently supports more than 700 federal laboratories with a total annual budget of over $23 billion, more than half of which is devoted to weapons development or other defense-related purposes, that is no small virtue.

References

Bilby, Kenneth. 1985. *The General: David Sarnoff and the Rise of the Communications Industry.* New York: Harper and Row.

Brock, Gerald W. 1982. *The Telecommunications Industry.* Cambridge, MA: Harvard University Press.

Ceruzzi, Paul. 1987. "An Unforeseen Revolution: Computers and Expectations, 1935–1985." In Joseph J. Corn, ed., *Imagining Tomorrow,* pp. 188–93. Cambridge, MA: The MIT Press.

David, Paul. 1990. "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox." *The American Economic Review, Papers and Proceedings,* vol. 80, no. 2 (May), pp. 355–61.

Devine, Warren, Jr. 1983. "From Shafts to Wires: Historical Perspectives on Electrification." *Journal of Economic History,* vol. 43, no. 2 (June), pp. 347–72.

Douglas, Susan. 1987. *Inventing American Broadcasting, 1899-1922.* Baltimore, MD: The Johns Hopkins University Press.

Du Boff, Richard. 1967. "The Introduction of Electric Power in American Manufacturing." *Economic History Review*, II. 20, December, pp. 509–18.
Fogel, Robert. 1964. *Railroads and American Economic Growth*. Baltimore, MD: The Johns Hopkins University Press.
Gelijns, Annetine. 1991. *Innovation in Clinical Practice*. Washington, DC: National Academy Press.
Maclaurin, W. Rupert. 1949. *Invention and Innovation in the Radio Industry*. New York: The Macmillan Company.
Martin, James. 1977. *Future Developments in Telecommunications*. Englewood-Cliffs, NJ: Prentice-Hall.
The New York Times. April 20, 1994, p. C1.
Rosenberg, Nathan. 1976a. "Technological Change in the Machine Tool Industry." Chapter 1 in Nathan Rosenberg, *Perspectives on Technology*. New York: Cambridge University Press.
———. 1976b. "Factors Affecting the Diffusion of Technology." Chapter 11 in Rosenberg, *Perspectives on Technology*. New York: Cambridge University Press.
Rosenbloom, Richard and Michael Cusumano. 1987. "Technological Pioneering and Competitive Advantage: The Birth of the VCR Industry." *California Management Review*, Summer, pp. 51–76.
Schurr, Sam et al. 1990. *Electricity in the American Economy*. New York: Greenwood Press.
Townes, Charles. 1968. "Quantum Mechanics and Surprise in the Development of Technology." *Science*, February 16.
U.S. Navy, Bureau of Ships. 1941. *Technical Bulletin* No. 2. January. As cited in James Martin, 1977. *Future Developments in Telecommunications*, p. 11. Englewood-Cliffs, NJ: Prentice-Hall.
Wriston, Walter B. 1992. *The Twilight of Sovereignty*. New York: Charles Scribner's Sons.

Discussion

Joel Mokyr*

Anyone can have fun collecting technological predictions that did not materialize. Time and again, learned scientists and skilled engineers have been caught predicting with Lord Kelvin that most things that can be invented already have been and that diminishing returns to innovation have set in. Inventions that we know now to have been of momentous importance were more often than not thought to be minor curiosa.[1] Nathan Rosenberg, in his stimulating and entertaining paper, thinks that in part it is a "failure of social imagination" that accounts for these erroneous predictions. I am not sure what a social imagination is, but clearly some writers, from Roger Bacon to Jules Verne to more contemporary science fiction writers, did not suffer from a lack of imagination. It is just that the worlds they imagined and the worlds that eventually materialized subsequently overlap very little.

If there is one technology that preoccupies and fascinates our current world, it must be decentralized information and communication. *That* was not what most people who imagined at all half a century ago foresaw. They believed that by this time space travel would be common—it is not. Many of them feared total control by a central government that had access to incredible technologies allowing them to manipulate people beyond anyone's wildest nightmares. Such totalitarianism has not come about. Kurt Vonnegut's nightmarish world in *Player Piano,* in which labor-saving technological change has made labor redundant, has not

*Robert H. Strotz Professor of Arts and Sciences and Professor of Economics and History, Northwestern University.

[1] One of the more entertaining examples was *The New York Times* prediction in 1939 that "Television will never be a serious competitor for radio, because people must sit and keep their eyes glued on a screen; the average American family hasn't time for it." Cited in *The Economist,* July 5th, 1996, p. 15.

arrived. In fact, we seem to be working harder than ever before, if Juliet Schor is to be believed. It is not, I submit, lack of imagination, but the simple fact that imagination had it mostly wrong. We are today the future of the 1950s, but—with some exceptions—we are not what futurologists of that time thought we were going to be. The Talmud had it right when it sighed that "since our second temple was destroyed, the art of prophesy was given to the fools."

Rosenberg distinguishes in his paper between "uncertainty" and "ignorance." Decision-makers do not have access to an even marginally informative probability distribution, he says, so ignorance is perhaps a more suitable concept than uncertainty. The uncertainty of technological change is one that society cannot hedge against and cannot diversify away (although, of course, individual firms engaged in R&D can). The question is why technological change is so difficult to predict and so difficult to understand, and Rosenberg gives some very good answers. My own answers may seem on the surface different from his, but at second glance will turn out to be more or less a reformulation rather than an alternative.

To start off, technological change involves two levels of uncertainty. One is the firm's *micro*problem: Will a particular line of research pay off? That question can be decomposed into a whole host of subquestions that compound each other: Can this technical problem be solved at all? Can *we* solve it? Can we do so before anyone else does? Will it sell, and at what price? The other level is the economy's *macro*problem: What kind of technological *regime* will emerge as dominant? Will it be using digital or analog computers? Western antibiotics or Chinese herbal treatments? Fixed-wing aircraft or dirigibles? Nuclear or fossil fuels? Boiled potatoes or oatmeal porridge? This is the kind of uncertainty that historians have to deal with when they wish to explain why a society's production techniques developed in one direction but not another, but it is also hugely relevant for decision-makers at the micro level.

Mutation and Natural Selection

The lack of predictability is not the curse of the economics of technological change alone. Evolutionary biology, too, is incapable of making accurate predictions.[2] Instead, biologists depend largely on the paleontologists and their fossil bones to tell them about the facts of evolution. No biologist has ever witnessed a speciation event. We know that speciation occurred, of course, from our past record, but speciation cannot be predicted, or even recognized when we see it. The concept only makes sense in view of the past. The reason why biologists cannot predict

[2] For more details on this analogy, see Mokyr (1991 and 1996).

is fairly obvious: Mutations occur at random. The raw materials from which natural selection has to choose have no systematic component. Direction is imparted to the system exclusively by the selection process, an ex post mechanism. Moreover, the selection mechanism picks new life forms by well-defined criteria, yet we rarely understand these criteria sufficiently beyond saying that they increase fitness, a rather circular argument.

Indeed, it is now well understood that many seemingly favorable mutations will disappear, for a variety of reasons. To be picked, a mutation may have to increase fitness, but the reverse clearly does *not* hold. In any case, we are a long shot away from making conditional predictions on what *will* happen. This does not mean we know nothing at all: We can make negative conditional predictions such as "If an insect weighing 300 pounds emerges by some implausible mutation, it will not survive." But narrowing the bands of the possible does not amount to prediction. The uncertainty thus comes in twice: We do not know what the supply of innovations will look like, nor do we know with certainty which ones will be picked for retention.

What does all this have to do with technological change? Social and cultural processes have increasingly been thought of in evolutionary terms (Campbell 1960; Cavalli-Sforza 1986). That is not to say that they resemble in all their details the mechanism we now think of as the neo-Darwinian orthodoxy in evolutionary biology. In the living world, because mutations occur as a result of random processes and are orthogonal to the "needs" of the organism, they are far more than likely to be either detrimental or neutral. Only in the rarest of cases will a mutation turn out to increase fitness. In other evolutionary processes, that is not necessarily so. Moreover, in biological evolution, parentage is either single- or biparental. In cultural and social processes, the analog of genomes can be acquired from many sources. Acquired characteristics are retained and passed on to other generations. And so on—the differences are quite substantial. All the same, many of us find it useful to think of technological change as an evolutionary process. Innovations occur and are passed through selective filters. Whether natural or not, the idea of *selection*, in one philosopher's catchy phrase, "Darwin's dangerous idea," is often thought to be central in explaining why things are what they are (Dennett 1995).

While cultural evolution (of which technological evolution is a special case) thus differs from what Darwinian dogma holds for biology, the idea of directed selection imposed upon an exogenous and stochastic supply of innovations seems to be a powerful notion. Predictability depends on the correlation between need and mutation. Rosenberg implies that while such a correlation may not be zero, it is not very high either. "Necessity" is neither a sufficient nor a necessary condition for an innovation to emerge. Needs remain unsatisfied despite frantic scrambling for a technological fix; at other times, as in the case of the Walkman cited by Rosenberg, the need may have emerged after the invention became feasible or turned out to be quite different than originally intended. What he is looking at is something that biologists have called "exaptation" following a term proposed by Gould and Vrba (1982): A trait may arise resulting from one set of selective pressures and then end up being used in an entirely different capacity.

Predicting the supply of innovation is thus an extremely risky venture. To make things worse, as noted, the selective filters are only very imperfectly understood. Why do some inventions succeed and others fail? It would be nice if we had a one-to-one mapping of "fitness" defined in some way (say, firm profitability) to the adoption of inventions. But we know better: Complementarities, frequency dependency, the fortuitous presence or absence of a crucial factor, the energy and single-mindedness of one individual (as in the case of Admiral Hyman Rickover and the heavy-water nuclear reactor), and other factors mean that contingency and luck will continue to play a role. Seventy percent of all new products that get through the first layer of filters and actually make it to the supermarket shelves disappear again in their first 12 months. If prediction were easy, such errors would not occur.

In short, then, two sources of uncertainty compound each other in technological history: one concerns which novelties emerged at all, the other which novelties that somehow emerged made it to the marketplace and survived. To repeat, this does not mean that we are *totally* ignorant, but that by and large the techniques that we end up using were not inevitable. Many artifacts and techniques actually in use are no more inexorable than the peacock or the platypus. Technological history, very much like natural history, is ridden by what we may call *bounded contingency*. The history of science is similarly the result of blind variation and selective retention, as a long series of distinguished historians of science from Donald Campbell to David Hull have been arguing. Some scientific advances are of course obvious, given what precedes them. The more we learn about how science developed, however, the more we start to understand how Kuhn's great paradigms often evolved as the result of political power plays and a directionality imparted upon science by the ways scientists made a living, not their internal logic. Science, in short, is no more predictable than technology.

The third leg of our triad of evolutionary processes is the changes in economic institutions. Douglass North, the guru of institutional analysis, has long called upon us to propose an evolutionary theory of institutions (North 1990). None has emerged so far. Perhaps this is because, unlike other cultural systems such as science and language, the intuition of what an innovation is, is less clearly defined. To be sure, institutional innovations, such as the emergence of modern stock markets, indentured servitude, or fee simple, have occurred, but institutional change seems to be less about innovation and selection than about adaptation and the emergence of certain conventions and coalitions that have a vague interpretation of Nash equilibria. All the same, it is fair to say that whatever process one envisages here, it would be foolhardy to construct models that predict what institutions are going to emerge in the future. Even if we could somehow specify the "demand" side, we do not always get the institutions we need and surely do not always need the institutions we have. Here too, history dictates what we can and cannot do. Present institutions are a Markov chain: the sum of all past changes plus an epsilon. Of course, sudden innovation is possible, and societies at times overthrow their institutional structure and pick another—but their choice is usually limited to what others have done before. Only a few times in history did a few societies have a true revolution (in that they set up a new set of institutions not previously tried by anyone else), usually with disastrous results.

The Coevolution of Technology, Science, and Institutions

Now that we have depicted history as these three evolutionary and unpredictable processes moving side by side, we can add another layer to Rosenberg's questions about predictability. The point is that technology, science, and institutions do not only evolve, they *co*-evolve. The path that technology can take is not only conditioned by its own past, its luck, and its selection mechanism, it is also conditioned by the unpredictable path of science and institutions. What is more, its evolution feeds back into the evolution of the other two. Such feedback could be positive, negative, or a mixture of the two. Many volumes have been written about how science and technology interact, and Rosenberg's ideas imply that even if the course of technology were entirely deterministic, we still could not predict its future. The same is true for institutions. Some institutions are conducive to technological change, such as choice in education, free labor markets, well-enforced property rights, intellectual tolerance, and political pluralism. Others are clearly detrimental, such as uncertain property rights, strong conservative labor unions, totalitarian government, and excessive conformism and deference to the achievements of past generations. Most have ambiguous effects, such as patent systems, religion, and democratic government. What is certain is that technology depends for its development on what is happening to institutions. If we cannot predict the one, we cannot predict the other.

One example of the coevolution of institutions and technology is especially relevant. I have repeatedly maintained that the success of technological progress depended not only on the marketplace and complementarities but also on the continuous struggle between those who want technological change and those who, for one reason or another, do not (Mokyr 1994). Inventions often encounter resistance, either from entrenched interests who stand to lose from the new technique or think they will, or from groups and individuals who, for one reason or another, do not approve of the invention. This is a source of uncertainty that Rosenberg does not mention, but not one to be glossed over. From the hapless Roman glassmaker who, according to Tacitus, claimed to have invented unbreakable glass and was assassinated for his troubles by the evil Emperor Tiberius, to a modern bioengineering company that has to cope with the likes of Jeremy Rifkin, animal rights advocates, and greedy lawyers anxious to skim off rents through product liability suits, innovators have had to deal with Luddites in all forms and shapes.

In addition, then, to the normal questions an inventor asks himself such as "Will it work?" and "Will I be the first one?" are the questions "What will my neighbors say?"and "Will the FDA approve it?" and "Could somebody sue me for product liability?" In some cases, such as the French inventor of the sewing machine, Barthélemy Thimonnier, and the English inventor of the flying shuttle, John Kay, the neighbors were unhappy to the point of burning down their workshops and forcing them into flight. Precisely because such resistance always takes the form of non-market mechanisms, the evolution of institutions friendly to new technology is crucial. It is this coevolution that is responsible for the relatively short duration of periods of rapid technological development. While technological and institutional development often aid and abet each other, leading to rapid technological change, institutions soon change and bring the process to a halt.

Rosenberg's policy recommendation, which is a call to "Let a hundred flowers bloom," seems to me very much in the traditions of evolutionary thinking, even if he does not acknowledge this: The process of innovation, he concludes wistfully, is inevitably inefficient and wasteful because of uncertainty. Yet this seems to me to miss the point that all evolutionary creativity is by necessity incredibly wasteful: Think of all the millions of mutations that go to waste before one is fixed. Think of the species that have gone extinct over the past 600 millions of years of multicellular life. In technological change, we may not even want to call this "wasteful" since the process itself cannot be made efficient. If it were, it would lose much or all of its creativity. Uncertainty is neither the *cause* of this inefficiency nor its *effect*. Rather, both are the results of the evolutionary dynamic at work here.

This coevolutionary dynamic is especially important if we are to understand the following sentences of Rosenberg's: "The existence and refinement of complementary technologies may exercise a coercive and conservative effect, forcing the novel technology to be placed inside the current 'system.' Alternatively, however, complementary technologies may be exactly what is necessary for the practical realization of an entirely new system." This seems to link in neatly with the current thinking on the dynamic behavior of evolutionary systems that are at the edge of chaos, to use Murray Gell-Mann's phrase (1994). They are neither in the purely conservative area in which all change is immediately absorbed in the existing system and any innovation is immediately frozen and absorbed in local equilibria, nor in the chaotic region in which minor change immediately causes a total disruption of everything and in which nothing ever returns to a predictable state. Instead, in this region of self-organizing complexity exist a finite number of equilibria in which the system can settle down, and while we cannot predict which one will be chosen, the choice is not totally random.

This is the kind of behavior that "new evolutionists" like Stuart Kauffman (1995) are trying to bring to bear on evolution. Economic history should take notice. Technological change may be one of these "supracritical" processes in history. For thousands of years, technological change occurred in a variety of societies, always to run out of steam and see the economy revert back to a steady state. The Industrial Revolution meant that all of a sudden technological change became the norm rather than the exception. At that stage, as Kauffman puts it, technology moved into the supracritical region, and "all bets are off." This is a pithy way of summarizing Rosenberg's main point.

Optimality and Adaptation

Finally, a debate among evolutionary biologists is highly relevant to Rosenberg's paper here, because it is mirrored in debates among social scientists interested in technological change. This is the debate between adaptationists and anti-adaptationists. The former basically maintain that evolution gets it right and that every trait that survives the harsh filters of natural selection has a purpose and a function. This is not quite equivalent to the "Panglossian" view that everything evolves into an optimum optimorum, but it does mean that there are no outcomes that are obvious and persistent errors. Gould and Lewontin's classic "Spandrels of San Marco" paper (1979) was a frontal assault on the "adaptationist program," as they called it.

In economic history we have similar debates: On the one hand we have Paul David (1986), Brian Arthur (1989), and others pointing to classic cases of lock-in as a result of path dependence, coordination failures, and externalities; on the other hand we have the standard neoclassical approach that maintains that if you observe a highly inefficient outcome such as the Qwerty keyboard you have not looked hard enough or you are overestimating the degree of inefficiency, as Liebowitz and Margolis (1990) have argued. The neoclassical view would not rule out contingency altogether, but it would confine it to more or less equivalent outcomes and deny that in the long run, chance and history could lead to persistent inefficient outcomes. If the internal combustion engine beat out electrical and steam cars, this view suggests, it is because it simply was better.

Rather than take a position on this debate—and Rosenberg must speak for himself—I think it is important to emphasize how and why technology here differs from living beings. A species that finds itself in drastically less favorable circumstances has to adapt in some way or it is likely to go extinct. It can adjust only if some members of the species have the genetic information that contains the raw materials necessary to adapt. If this is not the case, extinction is likely. In technological choice, adaptation is less constrained, since societies can adopt a completely different technology wholesale—at a cost. These costs are both private and social, involving at times quite radical adaptation. Think of societies that after millennia of use have to abandon the camel and the oxen and get used to the jeep and the tractor. What is interesting in this context is to investigate whether the private capturable benefits are sufficient to cover the costs. Without that, there may be another role for the government, but it is unclear whether the political environment will produce this correction. That, too, is a source of uncertainty.

References

Arthur, W. Brian. 1989. "Competing Technologies, Increasing Returns, and Lock-in by Historical Events." *Economic Journal*, vol. 99 (March), pp. 116–31.

Campbell, Donald T. 1960. "Blind Variation and Selective Retention in Creative Thought as in Other Knowledge Processes." In Gerald Radnitzky and W.W. Bartley III, eds., *Evolutionary Epistemology, Rationality, and the Sociology of Knowledge*, 1987, pp. 91–114. LaSalle, IL: Open Court.

Cavalli-Sforza, Luigi L. 1986. "Cultural Evolution." *American Zoologist*, vol. 26, pp. 845–55.

David, Paul A. 1986. "Understanding the Economics of QWERTY: The Necessity of History." In William N. Parker, ed., *Economic History and The Modern Economist*, pp. 30–49. Oxford: Basil Blackwell.

Dennett, Daniel C. 1995. *Darwin's Dangerous Idea: Evolution and the Meanings of Life*. New York: Simon and Schuster.

Gell-Mann, Murray. 1994. *The Quark and the Jaguar*. New York: W.H. Freeman.

Gould, Stephen Jay and Richard C. Lewontin. 1979. "The Spandrels of San Marco and the Panglossian Paradigm: A Critique of the Adaptationist Programme." *Proceedings of the Royal Society of London*, vol. 205, pp. 581–98.

Gould, Stephen Jay and Elisabeth S. Vrba. 1982. "Exaptation—A Missing Term in the Science of Form." *Paleobiology*, vol. 8, no. 1, pp. 4–15.

Kauffman, Stuart A. 1995. *At Home in the Universe: The Search for the Laws of Self-Organization and Complexity*. New York: Oxford University Press.

Liebowitz, S. J. and Stephen E. Margolis. 1990. "The Fable of the Keys." *Journal of Law and Economics*, vol. XXXIII, pp. 1–25.

Mokyr, Joel. 1991. "Evolutionary Biology, Technological Change, and Economic History." *Bulletin of Economic Research*, vol. 43, no. 2 (April), pp. 127–49.

———. 1994. "Progress and Inertia in Technological Change." In John James and Mark Thomas, eds., *Essays in Honor of R.M. Hartwell*, pp. 230–54. Chicago, IL: University of Chicago Press.

———. 1996. "Evolution and Technological Change: A New Metaphor for Economic History?" In Robert Fox, ed., *Technological Change*, 1996, pp. 63-83. London: Harwood Academic Publishers.

North, Douglass C. 1990. *Institutions, Institutional Change, and Economic Performance*. Cambridge: Cambridge University Press.

Discussion

Luc L.G. Soete*

As one has come to expect when reading Nathan Rosenberg's papers and books, his detailed analysis of the history of new technologies provides invaluable insights into the surprising ways the inventors themselves originally thought these innovations would affect particular activities. It is as if the level of inventors' creativity were somehow inversely related to their level of imagination with respect to possible applications of their inventions. In other words, a future remains for us economists and other creative social scientists!

Through the many cases detailed in Rosenberg's paper, one gets a strong impression of the predominance of widespread uncertainty in technological change. Five dimensions of uncertainty are emphasized:

1) the inherent difficulty of identifying uses for a new technology, given the often primitive condition in which it first appears;

2) the crucial dependence on improvements through complementary inventions, often in sectors where potential users are to be found;

3) the systemic features of such complementary improvements when society is confronted with entirely new technological systems;

4) the inventor's tendency to aim new technology at a narrow problem-solving task, thereby foreclosing possible applications in other, unanticipated contexts; and

5) the need for technological novelty to pass the test of cost-effectiveness: the economic test.

The author concludes, "New technologies are unrealized potentials that may take a very large number of eventual shapes." Hence, he

*Professor of International Economics, University of Limburg, the Netherlands, and Director of the Maastricht Economic Research Institute on Innovation and Technology.

strongly argues, governments should refrain from promoting any one technological alternative (the example of nuclear energy is given) or focusing on specific research support (he cites the example of the "War on Cancer"). Rather, "government policy ought to open many windows and provide the private sector with financial incentives to explore the technological landscape that can only be faintly discerned from those windows."

I am of course impressed with the variety of cases described by the author. While Rosenberg admits in his typical understated way that the evidence presented is "anecdotal," he nevertheless claims that the anecdotes have been so selected as to include many of the most important innovations of the twentieth century. I am, with my limited knowledge, tempted to agree, but first I would like to see a more rigorous and complete description of what could be reasonably called "major" (at what time, though, and on what basis?) new technologies over the past two centuries. Could it not be that precisely because of their unanticipated impacts, one tends to focus on just those technologies that turned out to have economic and social impacts so unexpected that they aroused the interest of historians? In other words, is the anecdotal evidence indeed just "anecdotal," that is, of little general value?

Second, is there not much more to be said about some kinds of sectoral or technological uncertainty associated with research? Surely, the uncertainties in new drug research are of a different nature from the uncertainties in designing a next generation of chips. In the first case, a much larger degree of uncertainty seems to be linked to the often "trial and error" nature of the research involved; in the second case, the research often appears to be progressing along a relatively straightforward engineering trajectory—with attempts at miniaturization, say, or use of alternative materials (see, for example, Moore's Law). Would it not then be reasonable to assume that the uncertainties in the latter case are much less uncertain, and much more the type detailed by experts in the field, like Robert Howe of IBM with respect to the future intelligent assistant, the potential impact of electronic networking on banking, and the commodification of financial services?

Third, and as indicated in the introduction to this session, is there not also evidence that these uncertainties might be different over time and might even display some cyclical characteristics, depending on the particular phase of the economy? Beginning some 15 years ago, Rosenberg and I have both taken part in a long-standing discussion with a German colleague, Gerhard Mensch, about the possible "clustering" of major new technologies in periods of depression (the Mensch claim) or recovery (our claim). Again, I would argue that beyond the particular long-term aspects of that debate, it has some significant features that could shed light on the way the aggregate performance of the economy might influence technological risks and uncertainties.

Let me elaborate somewhat on this point, in light of the paper by Dale Jorgenson as well. The "productivity paradox," highlighted in Rosenberg's paper, is illustrated for the G-7 countries in Figures 1 and 2 in a very approximate way. Productivity growth fell steadily in the G-7 countries over the past 30 years to about half the level of the late 1960s. (Productivity growth is measured in Figure 1 as GDP per man-hour for the aggregate G-7 and smoothed over 11-year averages; an even more rapidly declining picture would be obtained using growth in Total Factor Productivity.) This decline in productivity growth contrasts sharply with the increase in business expenditures on research and development over the 1970s and 1980s, as illustrated in Figure 2. It also contrasts sharply with the increase in the share of total private R&D funding in most OECD countries, which has also risen steeply over the same period.

In explaining this paradox, Nathan Rosenberg, along with many others including myself, emphasizes the numerous uncertainties and difficulties involved in identifying the efficient use of a set of relatively pervasive technologies, such as information and communication technologies. And again, while I am very sympathetic to this view and the optimism it gives rise to—"You ain't seen nothin' yet: The future is still going to bring us the major benefits of these new technologies!"—alternative explanations are possible, two of which I would like to highlight here.

The first one focuses on aggregate measurement issues and the likelihood that, increasingly, we are mismeasuring output in a large number of information goods and services. This explanation has been raised by many authors, and I will not elaborate on it here (for more detail, see Soete 1996). I would just insist, as has Nakamura (1995), that our failure to include "consumer surplus" in real output measures is likely to have led us to greatly underestimate output growth, by much more than we may have corrected through the use of techniques like hedonic pricing. My guesstimate is that Nakamura, with his assessment of a 2 to 3 percent per year overestimation of inflation, is probably nearer the mark than the Boskin report.[1] My own back-of-the-envelope calculation for Europe would suggest that 2 to 3 percent is a reasonable estimate of the overestimation bias there as well.

The second explanation focuses on the interaction between short-term macroeconomic policies and long-term incentives for investment in research and development, possibly the core question at this conference. The crucial question is whether one unexpected side effect of the monetary policies of the 1980s, which were aimed at reducing inflation,

---

[1] "Toward a More Accurate Measure of the Cost of Living." Interim Report to the Senate Finance Committee from the Advisory Committee to Study the Consumer Price Index, September 15, 1995.

Figure 1

GDP, Employment, and Productivity Trend Growth in G-7 Countries[a]

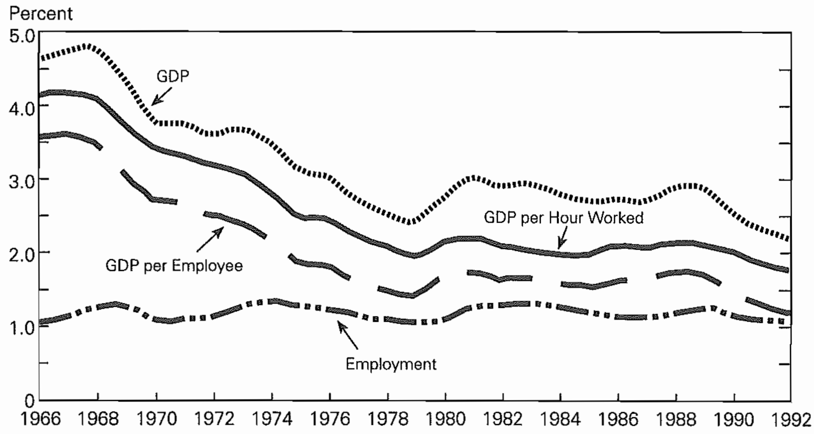

[a]Smoothed trend derived using centered 11-year moving averages.
Source: OECD 1996.

Figure 2

Investment Levels in G-7 Countries[a]

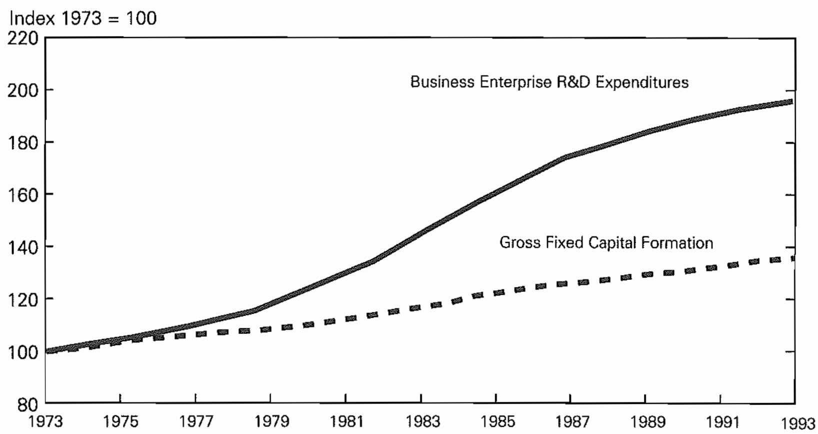

[a]11-year moving averages.
Source: OECD 1996.

Figure 3

Research Effort in the Business Sector

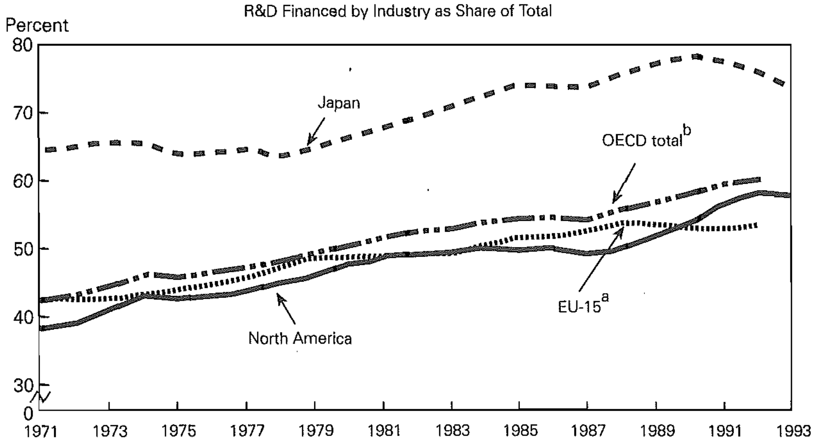

[a]EU-15 excludes Greece and Luxembourg in the 1971-80 period.
[b]OECD total includes EU-15, North America, Japan, Iceland, and Norway.
Source: OECD 1996.

has been a significant shift in the nature of research and development. In the 1980s and 1990s, real long-term interest rates turned positive and became, in the postwar context, extremely high. Among other things, the rise signaled the burden on capital markets posed by excessive government deficits in some of the major OECD countries, as well as by the growing capital needs of an expanding group of newly industrializing countries. As is illustrated in Figure 3, the 1980s were also a period of "crowding-in" of private R&D investment.

My claim would be that the short-term monetary policies of the 1980s led to a focus on short-term R&D, with a much stronger emphasis on product differentiation and relatively immediate private returns. In other words, short-term monetary policies could well have resulted in a pernicious decline in long-term potential growth, as high real interest rates shifted private firms' investment incentives to research of an immediate, short-term nature at the expense of longer-term, more uncertain efforts.

High real interest rates lead to an intensified focus on the present. Hence, more long-term, risky, and uncertain activities will often be eliminated from the research portfolio. Indirect evidence for such a trend, I would claim, can be found in the business literature on R&D and innovation management. This literature describes, in quite some detail, how R&D activities in many firms gradually became "streamlined" into business units' R&D centers. Strategic, "free" R&D had to become increasingly "legitimized" and was dramatically reduced over a very short period. As has been highlighted in consultancy reports, many large firms cut their strategic independent R&D dramatically: Hoechst, for example, from 75 percent to 25 percent of total R&D.

The result has been, as emphasized by authors in the innovation management area, that the R&D manager has become much more directly controlled by business unit managers, who are more aware of the immediate pressures for results. As Arnold (1992) put it: "Being close to the customer encourages incremental development and rarely inspires breakthroughs, simply because customers tend to have an evolutionary view of their need and rarely support a visionary spark." Similarly, a recent Arthur D. Little survey of European R&D managers points to the way "R&D functions are going through a quiet revolution, driven by intensifying competition and shorter product life cycles. They are becoming more closely linked to other parts of the business; researchers are becoming more aware of business economics and the needs of consumers." As one manager put it: "We have short-term profit and loss pressures which do not allow us to focus on long-term visions."[2] This case evidence fits well the aggregate trend and the resulting shift in the nature of R&D, described above.

To conclude, let me question the relevance of some of the historical analogies in this area. Can we really say, with anything more than faith, that as it took 40 years for electricity to produce efficiency benefits, the same must be true for current information and communication technologies? Surely the world has changed a great deal, and the technologies are by and large not comparable in their impact. The price decline linked to information processing reportedly already exceeds by a full percentage point the limited price effects of electricity. As Triplett (1994) has pointed out in a critique of an historical analogy made by Paul David (1990) between the computer and the dynamo, any simple diffusion model would tell you that the rate of diffusion of computer technology should be much more rapid than that of the dynamo. In other words, as time passes, productivity growth remains low; yet we are witnessing the introduction of ever more powerful information and communication equipment all around us, while technologists, economists, bankers, and policymakers herald the benefits of these new technologies. I am becoming more and more suspicious of explanations predicting the likely benefits to come on the basis of historical analogy. The future is not what it used to be, still less what historians today believe the past was like.

---

[2] See further Houlder, V. "Quiet Revolution." *Financial Times*, March 26, 1996, p. 10.

References

Arnold, R.H. 1992. "Pitfalls of Decentralisation, or Setting the Fox to Guard the Chicken Coop." *Research Technology Management*, vol. 35, no. 3, pp. 9–11.

David, Paul A. 1990. "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox." *The American Economic Review*, vol. 80, no. 2, pp. 355–61.

Nakamura, L.I. 1995. "Is U.S. Economic Performance Really That Bad?" Federal Reserve Bank of Philadelphia Working Paper No. 95-21, October.

Soete, L. 1996. "New Technologies and Measuring the Real Economy: The Challenges Ahead." Paper prepared for the ISTAT Conference on Economic and Social Challenges in the 21st Century. Bologna, Italy, February 5–7.

Triplett, Jack A. 1994. "Comments and Discussion on 'Computers and Output Growth Revisited: How Big Is the Puzzle?'" by Stephen D. Oliver and Daniel E. Sichel. *Brookings Papers on Economic Activity*, 2:1994, pp. 318–24.