

PAUL M. ROMER

University of California, Berkeley

National Bureau of Economic Research

Canadian Institute for Advanced Research

Implementing a National Technology Strategy with Self-Organizing Industry Investment Boards

THE MOST IMPORTANT LESSON from the study of research and development, economic growth, and the history of technology is that there are more ways to arrange the objects of the physical world than humans can possibly imagine. Ultimately, all increases in standards of living can be traced to discoveries of more valuable arrangements for the things in the earth's crust and atmosphere. The personal computer that I used to write this paper is made from almost exactly the same physical materials as the PC that I bought ten years ago—about thirty pounds of steel, copper, aluminum, plastic, and silicon, with bits of gold, iron oxide, and miscellaneous other elements mixed in. In my new PC these materials are arranged in a slightly different way that makes them about fifty times more useful than they were in the original configuration.

No amount of savings and investment, no policy of macroeconomic fine-tuning, no set of tax and spending incentives can generate sustained economic growth unless it is accompanied by the countless large and small discoveries that are required to create more value from a fixed set of natural resources. These discoveries are the product of a complicated set of market and nonmarket institutions that constitute what has been called a national innovation system. This paper considers both the economic

The author gratefully acknowledges comments from the discussants and the editors. This material is based on work supported by a National Science Foundation grant.

opportunities and the political risks inherent in attempts to strengthen this system. It proposes an institutional arrangement that could provide more financial support for innovative activity and direct it toward areas with large economic payoffs.

Even more impressive than the availability of a much faster PC to do word processing is the fact that my family drinks milk from a cow. If the milk cow did not exist, no one would ever believe that carbon, oxygen, hydrogen, and a few other types of atoms could possibly be assembled into a chemical refinery that automatically converts grass, water, and air into a nearly perfect liquid protein dietary supplement. This refinery operates with almost no human supervision, is mobile so it can search out its own inputs, can heal most mechanical failures, and can detect and neutralize any microscopic pathogens that enter the system. It even makes more than milk. It can build several replacement factories out of the same raw materials: grass, water, and air.

Compared with the cow, my PC has all the sophistication of a Tinkertoy. And if the particular arrangement of atoms that make a cow can come together through a blind process of mutation and selection—nature’s version of trial and error—imagine how many other ways atoms can be arranged that are as astonishing and valuable as a microprocessor and a cow. The fundamental challenge in economic growth is to find these new arrangements.

This optimistic potential for exploiting what Vannevar Bush called ‘‘the endless frontier’’ of scientific and technological opportunity is limited only by the difficulty of organizing collective action.¹ It takes collective action to encourage discovery and sustain the free flow of ideas, and the political mechanisms used to undertake collective action suffer from serious inherent defects.

But people create new institutions, just as they discover new technologies. Policy innovators discover better ways to undertake collective action, just as scientists, engineers, and product designers discover better ways to arrange physical objects. Vannevar Bush did more than write a report about the endless frontier of science. He midwived the birth of the National Science Foundation and peer-reviewed research grants for basic research at universities. He identified an important opportunity for collec-

1. Bush (1945).

tive action and then constructed an institutional arrangement to exploit this opportunity at minimal political cost.

Any contemporary discussion of a national technology strategy should be based on a balanced assessment of the potential benefits from collective action to spur growth and the risks inherent in undertaking collective action through the political process. Even die-hard free marketeers should concede that there is an opportunity for collective action to support some forms of basic research. On the other side, even the most ardent advocate of an activist government must admit that the nation's fifty-year experiment with an explicit industrial policy for the savings and loan industry shows how badly awry well-intentioned collective action can go when it is implemented through the existing political institutions.

Once the tension between the large opportunities for benefiting from collective action and limited political and institutional capabilities for undertaking effective collective action is recognized, it becomes clear that innovation in institutional mechanisms could be very valuable. Surely, the opportunity to experiment with new institutions is as great as the opportunity to experiment with new arrangements of physical objects. The standard dichotomy in economic policy debates—market exchange versus government intervention—does not capture the complexity of the kinds of social institutions already used to achieve common goals. It also fails to suggest the broad range of new institutional arrangements that could be tried. In its purest form the market is a mechanism for allowing independent action by all individuals, with no explicit coordination. The government is a mechanism for explicitly coordinating the actions of all people. Most economic activity is supported by institutional arrangements that are intermediate between these extremes.

After all, the modern capitalist economy is organized neither as a market nor as a government. It lets large numbers of people exploit the benefits of collective action through explicit, hierarchical coordination as members of a corporation, and in this sense large parts of a capitalist economy function like a government. But an economy with many firms subjects each to the discipline of competition, so in this sense the economy is organized more like a market. The diverse population of corporations is constantly changing, constantly finding new and different ways for large groups of people to work toward some collective end. Some corporations find new ways to structure themselves. Others die off and are replaced by

new and better arrangements. What a national innovation system needs is something like this kind of dynamic, something that lets us take advantage of mutually beneficial coordinated action and that uses the pressures of *competition and a market test to shut down ineffective institutional arrangements* and to reward promising ones.

To show how such a process might work, I describe a specific policy proposal conceived with this end in mind. Whatever the final judgment is on this particular proposal, the general analytical point will, I hope, be clear. To encourage technical change, it is not enough to call on business leaders to be more innovative. Policymakers must themselves be willing to experiment with new institutional arrangements.

Self-Organizing Industry Investment Boards

This proposal specifies a process, not a specific policy. To see how it would work, consider a hypothetical example. Suppose you run the Acme Widget company and I run Consolidated Widgets. You and I and many other widget makers come together and decide that there are industry-wide opportunities and problems that independent action by individual firms cannot address. You think, for example, that university professors could be doing useful research on questions about the basic principles of widget design and manufacturing if they just had the funds and the incentive to do so. I believe that the upstream industry that manufactures widget-making equipment could be designing more useful specialized equipment.

Under the proposal outlined in this paper, widget makers could take collective action on these matters by following the steps specified in a general piece of enabling legislation. They would start by petitioning the secretary of commerce, giving the argument for provision of an industry-specific public good. The secretary would then hold a hearing to certify that the proposal for collective action addresses a genuine public need. If the proposal passes this test, an election would be held in which all widget manufacturers vote whether to levy, say, a 1 percent tax on widget sales. If a large enough fraction of the industry (measured as a fraction of total sales, as a fraction of the total number of firms, or some combination of the two) votes in favor of this initiative, a tax

backed by the full force of law would be imposed on the entire industry. The proceeds from the tax would not, however, go to the government.

Part of the presentation to the secretary of commerce and then to the industry would outline plans for creating industry investment boards. You will take the initiative to organize an investment board that will fund university-based widget research. I will organize a board that supports the development of widget-making machinery in the upstream industry. Both of these boards will function as pass-throughs, accepting tax obligations from contributing firms and using these funds to support research in universities in the first case or development in upstream firms in the second. After the tax is passed, the two boards would solicit the tax funds from the firms in the industry. Suppose that your firm has sales of \$100 million a year, so it has to decide where to allocate its \$1 million in annual tax obligations. Because you are organizing the "University Research Board," you might decide to allocate all of the \$1 million to it. My firm has sales of \$200 million a year. I might decide to have my firm split its \$2 million tax obligation, giving \$1 million to the "Upstream Equipment Development Board" that I help organize and \$1 million to the University Research Board. Each year the leaders of other firms would have to decide how to allocate their tax obligations between these two boards and any others that might be created. The amount they have to contribute is fixed by their sales and the tax rate. They are free, nevertheless, to decide which board receives their contributions. If they do not approve of the boards that exist, they are free to start a new one.

Each industry board would have a board of directors answerable to its contributing firms and would operate as a private, nonprofit foundation. The boards would be limited by the general terms of the enabling legislation but would otherwise have wide latitude to make decisions without any direct oversight or second-guessing by the executive, legislative, or judicial branches of government. A general limitation would require each board to invest only in common property resources that benefit the entire industry. For example, all specialized equipment developed by the Upstream Equipment Development Board in collaboration with the upstream industry would be available for sale to all widget-making firms on equal terms. All research funded by the University Research Board would be freely distributed in the tradition of open university-based science. Thus, neither board could fund research or

equipment development that would be conducted in-house by a firm in the widget industry. All of the activities of the boards would have to be made public.

The enabling legislation should also specify that absolutely no tax funds could be used to support lobbying, public relations, or any kind of political activity. Also, no direct or indirect kickbacks or side payments to firms in the industry would be permitted. The tax rate used in each industry could vary, but a maximum tax, on the order of 2 percent of sales, would need to be specified in the enabling legislation. The legislation should also articulate the general principle that the tax should be a domestic consumption tax rather than a production tax. Units produced domestically for sale abroad would not be subject to the tax, but units produced abroad and sold domestically would. The legislation would also mandate equal treatment for all firms. Foreign firms would vote in the referendum, pay tax on their sales in the United States, allocate their tax obligations, and participate in the governance of the board or boards they support, in exactly the same fashion as domestic firms.

Suppose that other industry leaders come to believe that the most important problem facing the industry is an inadequate supply of post-graduate engineers with special training in widget-related design issues. Or suppose they decide that the most important investment would be to disseminate existing information about the principles of manufacturing to widget-making firms. Other firms might decide that junior college training of basic skills for current and future workers in the industry is the most important priority. An interested firm or group of firms could, at any time, petition the secretary of commerce, participate in public hearings, and then establish a Ph.D. Engineering Fellowship Board, an Extension and Diffusion Board, or a Worker Training Board. These new boards would compete with existing boards for the tax obligations of the different firms in the industry.

Over time, if contributing firms felt that the University Research Board was funding an incestuous network of scientists doing research of limited value to the industry, a competing research board could come into existence, or the funding firms could simply vote with their feet and take their tax dollars to other boards. In the worst case, if none of the boards were doing work that justified the cost, firms could use a periodically scheduled election to rescind the tax altogether.

The original proposal to the secretary of commerce would have to specify how the tax would be administered. Some boundary would have to be established, for example, between true widgets, which would be taxed, and near-widgets, which would not. An overall industry association, financed from an expressly limited share of total industry tax revenue, would have to be created to administer the tax system and perform the basic audit and information collection activities necessary to monitor compliance. Cases of explicit fraud or intentional noncompliance could be referred to an appropriate government agency with the power to impose fines, compel testimony, and, if necessary, undertake prosecution.

Precedents

This proposal may sound like political science fiction, but it is modeled on arrangements that already exist. The closest precedent is specified in the Agricultural Marketing Agreement Act of 1937, which outlines procedures for establishing arrangements called marketing orders.² Growers of a particular agricultural commodity can petition the secretary of agriculture to establish a market order on their behalf. Typically, the market order is approved if either two-thirds of the growers or growers representing two-thirds of output by volume support the proposal in the subsequent referendum. A new referendum is usually held every six years to gauge grower support. At any time, the secretary of agriculture can suspend any marketing order that is not operating in accordance with the aims of the enabling legislation.

Marketing orders designed to support research and development (R&D) and market promotion activities differ from the proposal in this paper only in that the funds raised by the tax automatically go to a single marketing board. There is no element of competition between boards and no free entry.

About three-quarters of the marketing orders in agriculture collect funds for research, development, and market promotion. (This is how the popular California raisin advertisements were financed, for example.) Marketing orders can also specify package and container regula-

2. Powers (1990).

tions, or size and grade regulations. These kinds of activities should be sharply distinguished from the notorious volume control regulations (that is, output restrictions) imposed in about half of all agricultural marketing orders. Quantity restrictions are clearly an activity that no economist would want to encourage, and they should not be tolerated, much less encouraged, in legislation that specifies the acceptable activities for the investment boards contemplated here.

This proposal has other, less obvious parallels with existing arrangements. For example, many (but not all) public electrical utilities contribute to the Electric Power Research Institute (EPRI), which finances research on a variety of industry-related matters. If local regulators let contributions to EPRI count as part of cost in figuring the utility's rate base, they are, in effect, using a government-sanctioned tax on electricity to support industry-related research. Just as in agriculture, however, there are no alternative research boards. Some observers maintain that EPRI may not be as effective in stimulating research as it once was or as it could be, perhaps because it faces no competition.

Before the breakup of the Bell System, Bell Labs was supported in a similar fashion. Each of the operating companies paid a few percent of total revenues to AT&T, which supported Bell Labs out of the proceeds. Again, to the extent that this contribution was built into the rate base that utility regulators allowed the operating companies, it amounted to a government-sanctioned tax used for industry-related research purposes. Because AT&T controlled the vast majority of telephone operating companies, free riding was not a problem.

Most observers judge Bell Labs to have been an extremely effective research organization. It made fundamental contributions to basic science. Information theory was created there. Radio astronomy was invented there. The background radiation that is the best evidence available for the "big bang" was discovered there. Bell Labs also produced high-quality scientific discoveries that have had enormous practical and commercial implications: the transistor, the laser, fiber optic transmission of information, and Unix, the first major computer operating system to run on computers made by many different manufacturers.

As evidence that firms can come together and take actions that are in the interest of the industry as a whole, one need look no further than the pharmaceutical industry. It recently persuaded the Food and Drug Administration (FDA) to raise the fees it levies when a company sub-

mits a drug for approval. The explicit understanding was that the FDA would use the additional revenue from this fee to hire more evaluators so that the agency could reduce the time it takes to reach a decision on drug approval.

The Economic Opportunities for Collective Action

It takes two arguments to make a case for this proposal. I must show why there are any important unexploited economic gains that require collective action. Then I must compare the potential gains with the political risks of implementing the system. This section outlines the theoretical argument for large potential gains from collective action. The next refers to evidence on their quantitative importance. Subsequent sections take up the political risks of this and other forms of technology policy.

If physical objects were the only economic goods, there would be little opportunity for collective action beyond the universally recognized need to establish a system of property rights. But as table 1 suggests, objects are not the only economic goods. They are not even the most important goods. The table presents a two-way classification of different kinds of economic goods. The left and right columns classify goods according to their costs of production. The horizontal dimension classifies them by the strength of the property rights available for each good.

The column on the left lists physical objects that are consumed directly or that provide services that are consumed. These kinds of things—land, fish, a worker's labor effort—are typically thought of as economic goods. The column on the right lists goods that can be represented as bit strings. They are all examples of information in the mathematical sense of the term. Whether it is literally a bit string (such as computer code or a digital musical recording) or is something represented in words or symbols that could be converted into bit strings (the design for a microprocessor, the operations manual for Wal-Mart stores, or the results of scientific investigation), any good in this column is a piece of pure information.

The technical terms from public finance for the types of goods in the two columns are rival goods and nonrival goods. The objects are rival

Table 1. Economic Attributes of Various Goods

<i>Degree of control (percent)</i>	<i>Rival goods (objects)</i>	<i>Nonrival goods (bit strings)</i>
100	Private goods: for example, a piece of unimproved land	An encoded satellite television broadcast
	_____	A digital music recording
	A car	The design for a microprocessor
		Computer code
	A worker's labor effort	The operations manual for Wal-Mart stores
		General principles of chemical engineering
		Principles behind window-based graphical user interfaces
		The do-loop in computer programming
	Fish in the sea	_____
	Clean air	
	Sterile insects used for pest control	Public goods: for example, basic research in physics
0		

goods because you and I are rivals for their use. You can eat the fish or I can, but not both of us. A bit string is a nonrival good because once it has been produced, we are not rivals for its use. I can listen to the musical recording or take advantage of the software code without in any way diminishing its usefulness to you or anyone else.

The most important result of the work on "endogenous growth" during the last ten years has been the renewed attention devoted to these attributes of ideas as economic goods. The first conclusion that emerges from the theory of growth was clear even in neoclassical models: it is the production of nonrival goods that makes growth possible. The second conclusion, which has emerged only after a great deal of work, is that the usual invisible hand result applies only to an artificial economy in which nonrival goods are provided exogenously by nature. In a real economy, an inherent, unavoidable conflict exists between the incentives necessary to encourage the production of these goods and the

incentives that lead to the optimal distribution of these goods, both to users and to the developers of other related nonrival goods. This means that private property rights and market exchange are not the perfect institutions for supporting growth. In fact, no simple description of the perfect institutional arrangement can exist. In any particular context, one must explicitly address the trade-offs both between the incentives for discovery and those for diffusion and between the limitations of market mechanisms and those of political mechanisms.

The two columns in the table correspond roughly to the distinction emphasized in the introduction between physical objects and intangible discoveries about new ways to arrange those objects. Farming, resource extraction, manufacturing, and distribution are all examples of actions that transform physical objects. For most people, these are the kinds of things that first come to mind when they think of economic activity. Discovery, in contrast, is how new, nonrival instructions are found for using rival goods more creatively. As the introduction suggests, discovery is where the real action is in economic life. Because nonrival goods are intangible, they are hard to measure. Moreover, they typically affect the economy in small increments that are hard to perceive from one year to the next. (The computer industry, where dramatic change can be seen in real time, is a notable exception.) But one needs to read only a little history to appreciate the profound cumulative effect of these nonrival goods—these discoveries—in almost all areas of economic activity. If the Earth were returned to the physical state that existed ten thousand years ago, wiping out all structures, physical capital, and civil engineering projects, but the total stock of accumulated knowledge were retained (in an exempted library where books and other records were kept), current standards of living would be recovered within a few generations. If the experiment were reversed, with the physical state of the world retained but the state of knowledge returned to what it was ten thousand years ago, our economic prospects would be much bleaker.

The sharp distinction that the table draws between rival and nonrival goods is rarely apparent in practice because real goods are almost always a mixture of the two. If you buy shrink-wrapped software, for example, you purchase a bundle that consists of rival objects (some floppy disks and a book) and a nonrival good (the legal right to use the bit string that encodes the computer program). The distinction between

an object such as a floppy disk and an intangible such as the computer code stored on the disk is important because it points to a fundamental difference in the associated costs of production. The cost of producing the computer code is virtually all fixed cost, or “first copy” cost. Firms spend millions (sometimes hundreds of millions) of dollars finding just the right bit string, but once they have it, they can replicate it at essentially zero cost. Once a nonrival good is produced, it is a good with no opportunity cost. The cost of the floppy disk, in contrast, is almost entirely the constant cost of producing additional disks; almost, but not quite, because, of course, the floppy disk itself reflects a sophisticated design that required an important fixed-cost investment. The floppy disk is not a pure rival good, but unimproved land is, because it has no underlying design cost.

In contrast to the horizontal distinction between rival and nonrival goods, the vertical dimension in the table reflects the more familiar issue of property rights, appropriability, excludability, or, simply, control. It is possible to control different kinds of goods by maintaining physical possession, perhaps with the assistance of the legal system, or by keeping some valuable piece of information secret. Land is placed relatively high up in its column because it is rarely stolen and because the cost of maintaining control over land is small compared with its market value. An automobile is lower down the column, because cars are more frequently stolen and the total resources spent by society on maintaining control over cars is higher. Goods that are both object-like (that is, rival) and over which almost perfect control can be maintained are called private goods. Land is a private good, and a car is close enough for most analytical purposes.

Farther down the rival goods column are examples of objects (or services from objects) for which property rights are weaker and control is less complete. When a firm hires a worker, it purchases labor services during certain periods of time. Because labor effort is difficult to observe, the firm does not always get the good for which it has paid. The existing legal system cannot enforce this kind of contract at all well. It would be absurd to propose that the firm go to court and sue for compensation whenever the employee leaves work early. Firms therefore find ways to execute transactions with their workers that are cheaper than writing explicit legal contracts that list all contingencies and then litigating every dispute. This is what the transactions cost theory of the

firm is all about.³ It emphasizes that firms are institutions that provide alternative systems for establishing property rights and enforcing contracts. These institutions make possible investments and gains from trade that would otherwise be impossible to exploit.

At the bottom of the column are examples such as fish in the sea or the sterile insects that are released in agricultural settings to control pests. Fish and insects are objects. They are rival goods. The insects can neutralize fertile pests in my valley or yours, but not both. They unambiguously belong in the left-hand column. Control over these goods, however, is especially weak. If one farmer paid for and released sterile insects, the benefits would spill over onto the farms of his neighbors. Much of the economic analysis of policy is framed in terms of these spillovers—these instances of incomplete property rights. Control can be weak, so spillovers can be present, both for rival and nonrival goods.

The right column lists examples of pieces of pure information with various different degrees of control. An encrypted satellite television broadcast, the kind used to distribute movie channels to cable television systems across the country, is a pure nonrival good with very strong property rights. Musical recordings, microprocessor designs, and computer code are examples of goods for which control is less than perfect but which nevertheless are supplied by commercial firms able to sell their goods in the market at a significant markup over marginal cost. The kind of knowledge possessed by the Wal-Mart corporation about managing retail stores is only partly controlled—other firms copy what it does. That knowledge nevertheless represents an asset in which a firm can invest and on which it can earn a sizable return.

Next come goods over which control is even weaker—these include goods such as the general principles that are the basis for chemical engineering, the insights behind the notion of a window-based graphical user interface for computer programs, or basic ideas such as the do-loop in computer programming. Finally, at the bottom, are goods, such as results from research in physics, that are nonrival and whose use is virtually impossible to control. Economists call these pure public goods. (The term is somewhat misleading because not all public goods are provided by the government—think of charitable contributions to

3. See, for example, Coase (1988); and Williamson (1975).

support public television. Moreover, not all goods provided by the government are what economists call public goods—think of the sterile insects.)

One of the most important insights in economics is that if people have strong control over ordinary objects (those in the upper left corner of the table where the private goods are) and if there are many potential buyers and sellers, decentralized exchange between self-interested traders leads to efficient outcomes. This is the lesson of *laissez-faire*, or the invisible hand. If control over objects is weak, outcomes may be inefficient—as in the proverbial tragedy of the commons. Everyone will be a free rider. These results are the basis for the strong, almost unthinking reaction among some economists that moving up the columns of the table and increasing the degree of control will always enhance efficiency. If others cannot be excluded from enjoying the benefits of some service, the government may collect taxes and pay for the service, as it does for the release of sterile insects, but this kind of collective action is used only as a last resort. Stronger property rights, this line of thinking suggests, would always be preferable.

This intuition is correct for rival goods but simply does not apply to nonrival goods, where strong property rights are inherently associated with monopoly power. If there are strong property rights, there cannot be many sellers. If firms that produce nonrival goods are to avoid large losses, these goods must sell for a price that is higher than marginal cost. Marginal cost on bit strings is zero, but the initial fixed costs of producing them can be very large.

To see why extremely strong property rights might be a problem, imagine that Bell Labs had been given a nonexpiring, ironclad patent on the discovery of the transistor. Or even worse, imagine that such a patent had gone to an organization such as IBM or General Motors. Think of how different the digital electronics and consumer electronics industries would be if every inventor who improved on the design of the transistor and every person who applied the transistor in a new setting had to negotiate with one of these large, bureaucratic organizations for permission to proceed. Not only would the prices have been higher, but the rate of discovery of all the inventions that reduced the cost of transistors would have been lower. These kinds of discoveries, which were made by many different individuals in many different firms,

have in just a few decades caused a millionfold reduction in the cost of producing a transistor.

A standard response to concern about strong patent rights is the claim that if GM had owned the patent on the transistor, it would have had an incentive to do everything right—to make all the right inventions and to agree to efficient contracts with outsiders. The relevant retort is that GM has had an incentive to find innovative ways to design and build high-quality automobiles and has not always succeeded in doing so. In the car business consumers are protected by the presence of a diverse set of competing manufacturers. But with an effective patent, all of the eggs would necessarily be placed in a single corporate basket.⁴

Or consider the computer software industry, an area of economic activity where producers are exquisitely aware of the need for some degree of intellectual property rights and monopoly power. Many thoughtful participants have also recognized that property rights can be too strong. Monopoly power can impose serious distortions and negotiation costs. Imagine, they say, if someone had been able to obtain a long-lived patent on the do-loop or the blinking cursor.

If people were pathologically honest and compulsively followed instructions, the economic problem of producing nonrival goods would be easy to solve. (This definitely is political science fiction.) In this kind of world, undertaking collective action would pose no problem. Everyone would diligently search out new opportunities for discoveries as they went about their other activities and would report all they learned to a central, coordinating agency. The government would direct a subset of these people to do the R&D necessary to take advantage of the most promising opportunities. The government could request that everyone else contribute a share of their income to the researchers who produced the desired nonrival goods—the software, the movies, the music, the books, the microprocessor designs, the innovative ways to organize a retail chain, the technological and scientific discoveries. These producers of nonrival goods could then give away all of the underlying discoveries at marginal cost, zero.

A world populated by real people instead of these science fiction

4. See Merges and Nelson (1992) for an extended evaluation of these kinds of costs from strong property rights.

automatons faces two distinct problems in providing nonrival goods: how to share costs and how to select the most promising opportunities for investment. Real people will choose to be free riders if they can. They will not share the fixed costs of goods that are freely disseminated if they do not have to. In addition, assembling all the information necessary to decide which of the extremely large number of possible nonrival goods to produce is difficult. The kind of calculation showing that there are many wonderful things to discover also implies that there are an almost infinite number of ways to waste effort on interesting but socially useless nonrival goods. Think about software. About $10^{1,000,000}$ different bit strings can fit on a 360K floppy disk. For comparison, a year has about 10^7 seconds, and about 10^{18} seconds have passed since the big bang. Out of $10^{1,000,000}$ possibilities only a very small fraction need to be useful for there to be many useful software programs still to be discovered. But so many other possibilities also are available that all human ingenuity for the rest of time could be devoted to producing useless computer code.

The government's powers of coercion make it uniquely capable of solving the cost-sharing problem. Unfortunately, these powers also make the government uniquely capable of wasting large amounts of resources on socially useless purposes. (Recall the experience in the savings and loan industry.) Markets, conversely, can solve the sharing problem only by introducing monopoly distortions, but they are better than governments at selecting the opportunities to pursue and avoiding wasteful spending. Because people operating in the market are motivated by the potential for profit, they seek out only those nonrival goods that have real value. The parallel or simultaneous search by large numbers of market participants can efficiently evaluate many possibilities. Bankruptcy constraints quickly cut off the flow of resources to projects that turn out to be unpromising.

Under the existing institutional arrangements for producing nonrival goods, one or the other of these extreme mechanisms is typically selected as being most appropriate for a given type of good. In the public good portion (the bottom) of the nonrival column, the government pays for basic research and gives away the results. This arrangement is chosen partly because dissemination of these goods is so important. (Think, for example, of the polio vaccine.) In addition, because of the efforts of Vannevar Bush and people like him, the institution of peer

review of competitive research grants is now available. It offers a reasonably good solution to the problem of selecting which projects to fund.

For nonrival goods at the top of the column, there is little prospect of setting up a government body that could make the right decisions about what to provide. No one would take seriously the suggestion that the government should extensively subsidize the production of popular music recordings, movies, or the design of new kinds of microprocessors so that these goods can be sold at marginal cost. Instead, society relies on market mechanisms to make those decisions and accepts the limits on dissemination and the monopoly distortions that the use of the market entails. The monopoly markups on compact disks, movies, and microprocessors cause relatively small welfare losses, and selecting which goods to produce poses an institutional design problem that dwarfs the problem faced in basic research.

The existing arrangement with government provision of basic research and market provision of final goods seems to work reasonably well for nonrival goods at the top and bottom of the column. It is the intermediate zone where the most important opportunities may now be missed. This region includes what Richard Nelson has called generic research. As he argues on the basis of case studies in different industries, this area may offer particularly large returns from investment in research.⁵ It includes goods such as the principles of chemical engineering, the insights behind the design of computer interfaces, and the fundamentals of program design. It is this region that my proposal tries to address through its mixture of government and private sector mechanisms. Without trying to identify in advance what these areas are and what the specific opportunity for collective action is, the proposal seeks to create a mechanism that combines the government's efficiency at solving free-rider problems with the market's effectiveness in selecting practical problems that offer the highest rates of return. Market participants can then make the right decisions about where the returns on investment are highest for the industry.

More is at stake here than just the rate at which knowledge is transferred from basic research to commercial application. These areas play a special role as the intermediaries between the basic research com-

5. Nelson (1983).

munity and the final users of technological knowledge in an economy. Better communication here is important not just for commercial outcomes, but for the vitality of the basic research endeavor. Without some point of contact with the practical opportunities and challenges of the world, basic science risks drifting into irrelevancy. The recent trend toward closer contacts between individual firms and universities avoids the problem, but firms pay for research only if they get proprietary control over the results. These arrangements therefore tend to undermine the traditions of open dissemination of ideas that have made our universities so successful. The proposal outlined here could provide a mechanism for connecting industry with universities, without jeopardizing the traditional role of the university.

Economic Magnitudes

Economists have uncovered a great deal of evidence suggesting the economic importance of nonrival goods. In his very useful survey of the econometric work on spillovers from R&D, Zvi Griliches describes one of the first attempts to compute a social rate of return on investments in R&D.⁶ T. W. Schultz computed the total resources saved by technological change in agriculture, compared the savings to total expenditure on R&D, and found a high ratio of benefit to cost.⁷ In his refinement of this calculation, Griliches himself found that the social rate of return on public investment in research on hybrid corn had a rate of return of about 40 percent, a number illustrative of the magnitude of returns on investment in R&D that has been found in many subsequent investigations.⁸ A sample of discoveries has been used to make this kind of calculation both in agriculture and in manufacturing.⁹ It has also been made in agriculture, in manufacturing, and at the national level by means of regression analysis using data on total factor productivity.

Three conclusions emerge from this large body of work. The first is that the social rate of return to investment in the broad class of nonrival goods is quite high, on the order of 30–50 percent. This level of return

6. Griliches (1992).

7. Schultz (1953).

8. Griliches (1958, p. 425).

9. Mansfield and others (1977).

confirms the claim made above that these are in some sense the most important kinds of investments that can be made. Second, the research demonstrates that the social rates of return are significantly higher than the private rates of return. This gap can arise because firms that produce new nonrival goods have only weak control or property rights over them. Alternatively, it can arise because property rights are strong; monopoly pricing then induces its own wedge between private and social values. The third conclusion is that economists cannot estimate the *ex post* rate of return in any one industry or area of economic activity with anything like the precision required if econometric estimates alone were to be used to make decisions about where to direct research dollars. This literature offers no support for the idea that academics or bureaucrats will be able to read the numbers and pick winners.

The calculation by Griliches suggests that the difference between private and social benefits from research is important for more than the microeconomic details. The macroeconomic effects of this difference can be quite large. Using the calculated social rates of return to investments in R&D, he can explain the majority of the total factor productivity growth at the national level as the result of measured spending on R&D. And this kind of calculation must lead to an underestimate of the importance of nonrival goods because it cannot capture the kind of innovation that led, for example, to discount retailing in the United States. Wal-Mart no doubt does not show up as a big player in the R&D statistics but has nevertheless helped transform an extremely important sector of the economy, significantly lowering costs in the retail sector. Griliches' calculation is also consistent with the recent cross-country regression estimates by Lichtenberg and by Coe and Helpman suggesting that social returns to R&D, measured at the national and international level, are still very high.¹⁰

Principles of Political Action

Having made the case that there are important opportunities for collective action to encourage the production of nonrival goods, especially in areas of practical importance, I now evaluate the potential costs of

10. Lichtenberg (1992); Coe and Helpman (1993).

trying to undertake collective action. In doing so several general principles suggest themselves. Other observers would no doubt add to the list and might change the emphasis, but the generalizations listed below should not be particularly controversial.

Reaction

Every policymaker operates in an environment characterized by competition among many different independent policymakers. A government of divided powers, such as ours, has a number of policymakers, and every government in the world must take account of the actions of other governments.

In this setting any change in policy can induce a reaction from other policymakers. A good working hypothesis is that policymakers operate in a strategic environment characterized by tit-for-tat and in which the details of what others are doing are sometimes hard to discern. This suggests a sequence of tests in evaluating any policy initiative:

- *Self-interest.* Would a policy be worth adopting if no other policymaker changed policy in response?
- *Reflection.* Would a policy be worth adopting if other policymakers responded by adopting the same policy?
- *Robust reflection.* Would a policy be worth adopting if other policymakers responded by adopting similar but more pernicious policies?

These tests reflect increasingly sophisticated views of strategic behavior. The self-interest test applies when there are no other players or in a one-time strategic interaction. The reflection test applies in cases of repeated interaction—what economists call a repeated game. The robust reflection test applies in cases of repeated strategic interaction if the actions or motivations of other players are difficult to observe—that is, in a repeated game with asymmetric information. In both of the reflection tests, the implicit model of the policy equilibrium is based on trigger strategies in which all parties defect as soon as one does.

One example that bears directly on technology policy and that suggests the importance of the robust reflection test can be cited.¹¹ In the

11. This account is based on Crease (1991) and Office of Technology Assessment (1991).

early 1980s George Keyworth, then the science advisor to President Reagan, invented the concept of a ‘‘Presidential Initiative’’ and simply put a \$140 million item in the Department of Energy budget to support funding for a National Center for Advanced Materials (NCAM), bypassing all of the usual mechanisms of peer review.

Suppose for the purposes of argument that this project represents a good investment in science and technology and thus passes the self-interest test. Suppose also that it passes the reflection test in a policy game between Congress and the executive branch. If both sides begin to use the special initiative process to fast-track high priority science and technology projects that are truly in the national interest, the net effect would be positive.

What Keyworth did not anticipate was that Congress would behave differently from the way he did because members of Congress face incentives that are different from his. They used the freedom from peer review to fund projects of dubious scientific value. At the same time that Congress rejected Keyworth’s attempt to bypass the traditional procedures with his presidential initiative (which was eventually funded after undergoing peer review), it approved two congressional initiatives that also bypassed all peer or agency review and that had no apparent justification in terms of national interests. For example, one initiative gave Columbia University \$24 million to renovate its chemistry building. (In a play on Keyworth’s words, this was referred to as the National Center for Chemical Research.)

These actions marked the end of the traditional consensus that basic research funding had to pass peer review and would not be eligible for congressional earmarking. Universities began to hire their own lobbying firms and compete actively for earmarked funding. Senators and representatives now compete to provide resources for their constituents. Since 1982 earmarked research funds have grown from \$9 million a year to \$470 million in 1991 and \$707 million in 1992.¹² For comparison, the annual budget for the entire National Science Foundation is about \$2 billion.

Almost everyone concedes that much of the congressionally earmarked funding represents pure pork-barrel spending. The academic community, the executive branch, and the defenders of traditional sci-

12. Savage (1992, p. 8, table 1).

ence in Congress have tried to put the genie back in the bottle but to no avail. For example, when the Bush administration tried to rescind funding for specific examples of pork-barrel science, Congress threatened an extended battle of tit-for-tat by canceling funding for thirty-one specific research grants for the National Science Foundation and the National Institutes of Health. The Clinton administration initially said the right things but has subsequently backed down and gone along with the intentions of Congress.

One can always argue that Congress would eventually have developed a high level of earmarked pork-barrel funding on its own, but the specific sequence of events and the testimony of participants suggests that Keyworth's initiative played a crucial role in the shift from the old equilibrium to the new one. If so, it did not pass the robust reflection test.

Delegation

The dissatisfaction with congressional earmarking as a mechanism for allocating funds to universities reflects a general principle. Successful decisionmaking on matters of science and technology policy requires that responsibility for specific spending decisions be delegated to a body that is not under the direct control of members of Congress. For example, delegation is the key element in the success of the peer review system. It is the mechanism whereby members of Congress commit to each other that they will refrain from pursuing special advantage for their constituents as long as everyone else refrains as well. The lack of delegation is at the heart of the failure of many science and technology projects that must, because of their large dollar cost, be directly funded by Congress. The summary by Cohen and Noll of the results from six case studies of large commercialization projects (the supersonic transport, satellite technology, the space shuttle, the Clinch River breeder reactor, the synfuels project, and the photovoltaics commercialization program) is illustrative. "The overriding lesson from the case studies is that the goal of economic efficiency—to cure market failures in privately sponsored commercial innovation—is so severely constrained by political forces that an effective, coherent national commercial R&D program has never been put in place," they wrote.¹³

13. Cohen and Noll (1991, p. 378).

The problem with direct political control of large discretionary spending projects is one of incentives, not ethics. Members of Congress are rewarded in large part for providing services and benefits to constituents. This creates strong incentives to select projects for reasons that do not have much to do with economic efficiency or even institutional missions. For example, the only possible explanation for the fact that the National Aeronautics and Space Administration (NASA) has been authorized to build the Christopher Columbus Center of Marine Research and Exploration in Maryland, a project the agency did not request, is that the senator who chairs the panel that writes the NASA appropriation is from Maryland. Even if the ultimate goal of this or any other senator is to do the very best job possible of promoting national science and technology objectives, the first priority is to build a strong local power base to ensure reelection.

From an institutional point of view, the problem is further compounded, because what Congress can delegate, it can also take back. A consensus that everyone should refrain from seeking narrow advantage in some area of government activity can easily break down if the opportunity presents itself and the temptation becomes too large. Some proponents of direct government funding for specific kinds of commercial innovation point to the former Defense Advanced Research Projects Agency (DARPA) as an agency comparable to the National Science Foundation, able to make independent decisions free from congressional interference. This may have been true in the past, when its activities were clothed in secrecy and were not widely understood. But lawmakers from Michigan inserted \$25 million into the fiscal 1993 defense budget to fund purchases by what is now called ARPA of flat panel display screens. (To emphasize the increased importance that this agency is to place on technologies that have both civilian and military applications, the new administration has taken “defense” out of the name.) It will come as no surprise that the legislation was written so that the company best positioned to fill the request happened to be located in Michigan.¹⁴ Officials at ARPA reportedly disapproved, but they know where their funding comes from. They wrote a request for proposals that favored the Michigan firm and are now negotiating a final purchase contract.

14. Carey (1993).

Inertia and Competition

A third general principle is that large organizations are very difficult to change. They seem to adapt to changing circumstances only when confronted with serious competitive threats. This generalization seems to apply equally well to private and public sector organizations. Large companies such as Ford and Chrysler, which are now two of the low-cost, high-quality automobile manufacturers in the United States, began the major reorganizations that transformed their operations only after being faced with the serious prospect of bankruptcy. (Some economists point to the government's intervention to save Chrysler as a success. Perhaps, however, the government should have demonstrated its willingness to let a big-three automaker fail. If it had, Ford and especially GM might have begun to change much sooner.) Similarly, while the rest of the Soviet economy fell further and further behind Western levels of output and technology, the Soviet arms industry—propelled by the intense pressures generated by military competition—managed to manufacture fighter aircraft that were roughly on a par with those made in the United States.

This suggests a role for competition that is quite different from the one economists teach to students. Competition in this sense has nothing to do with price-taking and tangency conditions. Instead, it affects the power that the leaders of a large hierarchical organization have to overcome the incentive and monitoring problems that tend to plague these institutions under normal circumstances.

Reaction, delegation, and organizational inertia are general considerations that policy analysis in any country of the world should address. The next two issues are somewhat more specific to the current political and institutional context in the United States.

Pay as You Go

In the current budget climate in the United States, no large new spending program can be undertaken unless it includes a funding source. The magnitude of the change in spending priorities that an aggressive technology policy might contemplate makes the importance of this observation clear. The substantial difference between the social and private rates of return to investment in nonrival goods and their relatively small share of total investment suggests that it would be

reasonable to consider an increase in annual spending on R&D of as much as 1 percent of gross domestic product, or roughly \$60 billion a year. This is not an impossible amount to contemplate in a \$6 trillion economy, but it is a very large amount to raise in new taxes in the midst of the current budget climate in Washington.

Divided Government

The American system of government was intentionally designed to limit the ability of any one part to take decisive, unilateral action. Subsequent developments, such as the recent erosion of the seniority system and party discipline in Congress and the expansion of the powers of the judiciary have probably pushed the system even further in this direction. The government is not an actor. It is a game with a very large number of players, many of whom can veto, or at least delay, any proposal. The equilibria that result are quite different from those that emerge, for example, under a parliamentary system.

To see the relevance of divided government for technology policy, consider two examples. First, ARPA (and before it, DARPA) has explicitly encouraged technological development of massively parallel supercomputers. It picked particular firms that it aided during the development of prototype machines. Firms that were not favored have objected to the perceived favoritism and have managed to instigate, no doubt through their local congressional representatives, a General Accounting Office investigation of ARPA that was critical of its supposed favoritism.

Set aside the question of whether ARPA is right that massively parallel computing is an important area of technology. (It almost surely is.) Also set aside the more problematic question of whether ARPA is able to pick the most promising firms in this area. The relevant point is that any agency directly answerable to Congress is increasingly unable to exercise any independent judgment about important technical issues. Soon, the selection of technology projects at ARPA will look like the rest of military procurement.

Military procurement is, of course, a familiar whipping boy. The Air Force, for example, has been trying to buy a large quantity of desktop computers for years now under its new "expedited" bidding process. The winner has been announced twice, only to be overruled when losing

bidders protested to the General Services Administration board of contract appeals. As James Q. Wilson has emphasized, the inflexible, rule-driven, excessively bureaucratic nature of procurement is an inevitable outcome of a process in which a losing party has access to multiple bodies—members of Congress, other agencies, or the courts—to question the integrity and professionalism of the person who made the decision to give the bid to someone else.¹⁵

Evaluating the Self-Organizing Industry Boards

The economic case for creating the industry investment boards has already been made. The opportunity for collective action is a logical consequence of the existence of nonrival goods. (So, for that matter, is the potential for economic growth.) The private sector-government hybrid proposed here would use the tax system to share the cost of producing new nonrival goods, and it would use market forces to solve the selection problem of deciding how the funds could best be invested. Private, for-profit firms would make the decisions about the relative importance of various industry-specific public goods and whether these interests are sufficient to justify any funds at all. In addition, both the industry boards and the organizations, such as research universities, research labs, and training schools, that they might support would be forced to compete for funds. Thus, both private sector incentives and competitive pressures can be brought to bear on the selection process.

If it succeeds, this mechanism could support investments that span the gap between the most general forms of basic research and the product development activities of individual firms. During the past few decades new discoveries in this gap have been funded only if they bear on the health care or military missions of the federal government. Because other practical areas have been neglected and because support for health- and defense-related research is likely to contract in the coming years, it is particularly important to find a new way to select the practical problems that motivate good fundamental research and create economic value. A new source of funds to support work on these problems must also be found.

15. Wilson (1989).

Another, more subtle advantage of this kind of proposal is that it addresses broader issues than those typically considered in the technology policy debate. As Henry Ergas has emphasized, explicit government policies affecting procurement, subsidies for firms, diffusion, or the educational system take place within a broader economic and institutional context that crucially determines their effectiveness.¹⁶ Implementation of this proposal can empower firms in various industries to change some of the basic features of this larger context. They can change, for example, the patterns of interaction between university researchers and those in industry. They can change the system of training at the postgraduate or vocational level. They can create new institutions for setting standards. In so doing, they can enhance the effectiveness of the policies that governments now take (too often with little success) to encourage the commercial application of new technology.

Some of the details of the proposal could be debated but are of secondary importance. For example, the primary effects of selecting a specific tax will come from investing the revenue in an area with large disparities between private and social returns. The distortions associated with different kinds of taxes will be second-order small. The decisive issue in selecting the tax will come not from the conventional analysis of elasticities and deadweight losses, but from the costs of collection and enforcement. Industry participants would have the right incentives to minimize these costs and should therefore be given wide latitude to propose their own enforcement mechanisms.

The political case for this particular proposal is somewhat more complicated. This system was explicitly designed to survive the robust reflection test. If every country decided to compete internationally by creating private sector initiatives that support industry-specific versions of Bell Labs, or open versions of SEMATECH, or new schools of biotechnical engineering, everyone would benefit from the discoveries that result. The federal government would certainly want to insist that other countries offer the same treatment to U.S. firms that they offer to domestic firms. If such a system were adopted in another country, the U.S. government would also insist that the spending must be transparent and that it not be used to offer continuing subsidies to domestic firms in an industry. If differential treatment or direct subsidies were possible,

16. Ergas (1987).

there is a real risk that even if the United States did not abuse this kind of arrangement, other countries could and would.

This proposal is also designed to ensure an adequate degree of delegation. The only effective way to do this is to make sure that the decisionmakers can operate without having to go through the congressional appropriations process each year. (This is why the Federal Reserve Board is independent in a way that no other quasi-governmental entity is.) Absent this freedom, the kind of interference that is beginning to occur at ARPA is sure to arise here as well. Finally, the provisions specifying free entry of new industry boards are designed to ensure that the inertia that typically saps the effectiveness of any complacent organization will not affect prospects for an entire industry—or for the nation as a whole.

In more parochial terms, the proposal is obviously designed to be self-financing and is tailored to keep any new financing arranged under this proposal isolated from the ongoing and protracted debates about federal budget priorities and deficit policy. As the example of the FDA fees suggests, it is not entirely unrealistic to believe that firms in some industries would be willing to back the required taxes if they could control how the proceeds are used. Finally, this kind of system, unlike any government agency, would be capable of decisive action because it would be free from interference from the legislative and executive branches of government.

One of the most serious (and certainly the most frequent) complaint about this proposal is that it may not work for all industries, perhaps not even for very many of them. Compared with a multibillion dollar program that puts control of innovation into the hands of some technology bureaucrats, this proposal is indeed less aggressive. But if one believes that the risks from careless intervention by the government are potentially very large, this built-in conservatism is an advantage, not a weakness. The imperative in policy design, as in medicine, should be to do no harm. In its current form, this proposal will be exploited only in industries whose participants perceive benefits that are large enough to justify the costs to them of the tax (which may in fact be small) and the costs of setting up and administering the system (which may be large). Where it is implemented, society will very likely derive important gains. Where it is not, the case for other, stronger measures must

be viewed as open in principle, but with a very strong burden of proof on those who call for more explicit government intervention. Surely, the right way to proceed in formulating policy is to try the less risky, less intrusive measures first and save the measures favored by advocates of aggressive government intervention for later.

If this criticism is right, this mechanism will be implemented in few places. If so, there is less to be gained by creating this process, but little risk in trying. And the criticism may well not be correct. For many areas the opportunities seem to be large, and the mechanics of collecting the tax straightforward. For example, the telephone and cable TV industries could agree to a common, 1 percent revenue surcharge that would be used to support research and training on new digital communications and the interface between computing and imaging.¹⁷ In effect, these industries could recreate the source of funds that once supported the rich interaction between practical problem-solving and fundamental research that took place at Bell Labs. Given the increasing likelihood that the “Baby Bells” will end up directly competing with each other (for example, through links with cable firms that operate in the service area of another phone company), the existing arrangements that support BellCore may erode. In any case, BellCore, the part of the original Bell Labs now controlled by the telephone operating companies, probably could benefit from being exposed to a more competitive research environment.

One could imagine an industry initiative by semiconductor manufacturers, with a tax of, say, 20 cents per million transistors on all domestic chip sales. If a firm in the semiconductor manufacturing industry believes that the SEMATECH consortium deserves support, it could use

17. After this paper was written, I learned of a research organization called CableLabs in Boulder, Colorado, that collects a voluntary charge of 2 cents per subscriber per month from the operators of cable television systems. Contributing firms cover 85 percent of the subscribers in the United States and about 70 percent of subscribers in Canada. CableLabs employs a small number of scientists and contracts out much of its research to universities, public and private research laboratories, and some equipment suppliers to the industry. It is currently working on problems such as digital compression and digital transmission of cable signals.

The fact that an organization very close to the one described here could come into existence even without the support of legislation that solves free-rider problems suggests the existence of a much larger unmet demand for the kind of collective industry-specific investment activity than skeptics have realized.

its funds to ensure support once ARPA withdraws its funding, as it is now scheduled to do. SEMATECH could compete with universities for funding from an industry-controlled research board, provided it operated according to the open disclosure and equal access rules required under the new system. If the research board did not believe that SEMATECH was a good investment, it could use its funds to support basic research at universities or the national labs. Other boards could support graduate training fellowships, worker training, or whatever other need they perceived to be most pressing.

One could also imagine initiatives in machine tools or in automobile assembly and design. More prosaic economic areas such as construction might also be covered, perhaps in subsegments such as single-family home construction, an area that seems to have been conspicuously lacking in technological progress.¹⁸ A new initiative for electrical power utilities might compete with EPRI in areas such as electricity storage. Biotech firms could use their funds to create a school of biotechnical engineering that could do for them what chemical engineering at MIT did for the petroleum and chemical industries in the United States.¹⁹ Software firms could tax shrink-wrapped software to help pay for basic research in computer science or to establish better curriculum standards, testing, and teaching at the junior college and undergraduate level. Software firms might also benefit from the creation of a software engineering discipline separate from the pure research activities of existing computer science programs. These schools could support the systematization of private sector knowledge about principles of software engineering and could train skilled professionals for software production rather than for university research, just as chemical engineering serves these functions for the petroleum and chemical industries. This kind of change in the institutional infrastructure could prove to be far more important for an industry than any particular government program of subsidies or attempts to pick winners.

In any industry complicated details of administration and tax collec-

18. See Nelson (1983) for a discussion of previous attempts to transfer the successes of the extension service in agriculture to the homebuilding industry. At least part of the problem in previous initiatives has been the attempt to implement the program from the top down, rather than meeting needs that were perceived by industry participants.

19. Rosenberg and Nelson (1993).

tion would have to be worked out. For example, if a tax of 20 cents per million transistors were levied on semiconductor chips, some decision would have to be made concerning the level at which tax obligations would be calculated and monitored. Typically, firms try to avoid responsibility for taxes, but under this proposal, control of how the funds are spent comes along with responsibility for collecting the tax. A tax per transistor could be levied when a semiconductor manufacturer such as Intel sells one of its chips. Intel would be responsible for reporting its total sales of transistors and would be able to decide which industry research board would get its taxes. When Toshiba exports memory chips to the United States, it could be responsible for reporting sales and could allocate its tax obligations among different boards. Alternatively, a firm such as Dell Computers that buys semiconductor chips from Toshiba would be responsible for reporting its purchases of transistors not already covered by a tax payment and could decide how these tax obligations would be allocated. One could even leave the choice of who pays the tax and controls the revenue up to the firms involved. Toshiba could sell some chips on which the tax has already been paid if it wanted to control some of the tax revenue. It could also sell some chips on which taxes have not been paid if Dell or another purchaser wanted to pay the tax and control the proceeds. All that would be required is that someone pay the tax before the goods move into the hands of the final purchaser.

A little thought would be needed to resolve the practical details, such as implementing a system for keeping track of the chips on which tax payments have already been made and those on which they have not. This problem is the kind that a government agency might find very difficult to solve or might solve by imposing large costs on the private firms that have to comply with its rules. But private firms routinely solve all kinds of difficult logistical, contractual, and monitoring problems. Dairy cooperatives make sure that fresh milk is always on the shelf in the grocery store. Railroads keep track of freight cars that are shuttled among different trains operated by different companies. Employers make sure that workers actually do their jobs in the absence of effective legal enforcement of employment contracts. Policy analysts should not underestimate the ingenuity of private firms in solving monitoring, enforcement, and tax collection problems in their industry if

they perceive it to be in their interest to do so. And if firms do not perceive this program to be in their interest, it will not even be attempted.

Perhaps the biggest risk inherent in this proposal is that existing firms in an industry might use it to protect themselves from new competition. For this to take place, some degree of collusion between existing firms and the government would almost certainly be necessary. The requirements for openness and the oversight powers of the secretary of commerce would make it very difficult for existing firms to limit competition without at least tacit government approval. Unfortunately, cases in which the government would collude with the threatened firms in an industry are not hard to imagine. For example, the United Steelworkers have proposed a surcharge of \$5 a ton on all imported and domestic steel to cover unfunded pension and health care obligations at large steel firms. This is a clear attempt to tax the new, innovative, nonunion minimill steel producers and subsidize the existing firms and their workers. The United Auto Workers are also considering this kind of arrangement for the automobile industry. Because the government, through its Pension Benefit Guarantee Corporation, is faced with large potential losses on the pension obligations of the largest steel and auto firms, it might be willing to go along with this kind of scheme.

Nothing prevents Congress from passing this kind of surcharge as a separate measure, but it could not be imposed surreptitiously under the structure of the industry investment boards. The enabling legislation would specifically prohibit any kind of spending activities by boards that would shift costs from or transfer resources among existing firms. The ultimate protection, however, comes from the freedom that firms in an industry have to create alternative boards and to control how funds are spent. Even if the established steel producers, because of their size, could impose an industry-wide tax that minimill operators do not want, the established producers have no way to coerce the minimill operators to fund an Industry Pension Board. The minimill operators could merely establish a separate board devoted to research on new technologies for minimills and refuse to contribute to any board dominated by the old, integrated producers.

The freedom to create new boards and to choose among competing boards is so important that any proposal that did not include these

provisions would have an entirely different character. Without these protections, the political risks would loom very large compared with the potential economic benefits.

Alternative Technology Policy Proposals

The criteria listed above can be used to evaluate other technology policy proposals.

Increased Funding for Basic Research at Universities

Funding for basic university research has until recently received broad support from economists and politicians, partly because the economic advantages of free dissemination of basic scientific discoveries are so obvious. In addition, the delegation problem has largely been solved. Peer review delegates resources to the scientific community, which decides what issues to pursue within broad guidelines established by the funding agencies. (But as noted above, this delegation function is increasingly threatened by direct congressional earmarks that bypass peer review.)

More and more policymakers, however, are realizing that the existing arrangement solves only one part of the broader selection problem of deciding which areas deserve additional research support. Spending is concentrated on pure scientific research and on practical problems in the areas of health and defense, but too few resources are devoted to areas of basic research and training that are motivated by the commercial opportunities faced by private firms. Because policymakers are beginning to emphasize the commercial relevance of research and because budgetary and geopolitical changes may reduce mission-oriented basic research in health and defense, the amount of support universities receive for basic research may soon decline substantially. From a practical point of view, the relevant question is not how much government support for basic science at universities can be increased, but whether the current level can be preserved.

To understand how a system of industry investment boards could protect the existing level of support for pure research at universities

and at the same time address the reasonable concerns of policymakers for a commercial payoff from some kinds of research, it is useful to look at the history of universities in the United States. As Nathan Rosenberg and Richard Nelson demonstrate, the university system in the United States has always paid a great deal of attention to practical problems and has contributed enormously to their resolution.²⁰ But the nature of those practical problems has changed dramatically over the years.

Before World War II universities primarily addressed practical problems faced by commercial interests. The land grant universities were devoted to agriculture and “the mechanic arts.” Farming, railroad transportation, electrical generation and transmission, and chemical processing all required their own specific kinds of technical research and training. Europeans sneered at the vocational orientation of higher education in the United States during the nineteenth century and the first part of the twentieth century, but Americans derived important national economic benefits from a responsive university system. To illustrate that responsiveness, Rosenberg and Nelson describe the emergence of programs of electrical engineering. In 1882—the same year that Edison opened the first electrical station in New York—MIT introduced its first electrical engineering courses. Cornell followed in 1883. By the 1890s schools of electrical engineering were the primary suppliers of skilled professionals for the electrical industry, having supplanted unsuccessful attempts by firms to train engineers in house.

This experience was not unique. Universities trained people for work in many areas of commercial importance and developed bodies of knowledge that could be used there. They did not do this, one can be sure, entirely out of patriotism or dedication to national economic success. Universities were attentive to private sector opportunities because the private sector paid the bills.

During and immediately after World War II, the federal government became by far the largest patron of research and advanced training at universities. This shift in funding accompanied a shift toward solving practical problems in health and defense and pursuing more abstract problems. The Department of Defense and departments with related

20. Rosenberg and Nelson (1993).

missions (Energy and NASA) provide the majority of all research support in the engineering disciplines. The total expenditure at the National Institutes of Health is five or six times larger than the budget at the National Science Foundation.

There is a continuing debate about whether this postwar shift in emphasis has been beneficial for the business sector. Many observers say that the commercial spin-offs from military research have decreased in recent years. The typical interpretation of this finding is that civilian and military technologies are diverging for exogenous technological reasons. An alternative explanation is that the initial surge in funding from the government in the immediate postwar years stimulated a system that for historical reasons was still attentive to commercial applications and opportunities. As the system evolved in the 1960s and 1970s, however, its connections to the civilian business sector atrophied. The increasing divergence may therefore reflect the delayed effects of incentives and funding rather than exogenous technological developments.

However this debate is resolved, the end of the Cold War clearly means that defense-related support for research will be substantially reduced in the next decade. Moreover, the enthusiastic support for technological advance in medicine and health may increasingly come under attack as the nation tries to come to terms with the rapidly escalating cost of medical care. Some policy analysts are already suggesting that because new health technologies are not being priced and allocated correctly, slowing the rate of technological advance would be the best way to control health care expenditures.

Stagnation, or even real cuts in government funding for university research, is thus a distinct possibility. Moreover, demands will grow for more direct commercial payoffs from research. This pressure is reflected, for example, in the controversial planning process that Director Bernadette Healy tried to impose on biological research funded by the National Institutes of Health in 1992. Responding to congressional pressure, she tried to adjust the decision process and priorities of the traditional peer-reviewed, basic research program to give greater emphasis to commercial objectives. In the current institutional environment, this pressure could undermine the strengths of the existing basic research system without generating the hoped-for commercial benefits.

The National Institutes of Health, the National Science Foundation, and the Department of Defense are unlikely to be the best institutions for assessing commercial opportunities.

A more productive response to politicians and voters who ask for research and training that have closer economic payoffs would be to agree to a division of labor. Basic science agencies such as the National Institutes of Health and the National Science Foundation could continue to fund the most adventurous and forward-looking research at their existing levels, and an independent source of funds could be created for the commercially relevant research that should be under the control of people in the private sector who are knowledgeable about the opportunities. These funds could fill in part of the gap that would open if mission-oriented funding in defense and health is cut back. Such a dual system could both restore the strengths of the pre-World War II connections between universities and the private sector and preserve the advantages of the more abstract basic research system that developed after the war. It might also provide more total funding than is currently available. The system of industry investment boards would establish this more balanced arrangement.

It is important to recognize (as Rosenberg and Nelson do) that an increased emphasis on practical problems is completely consistent with a division of labor in which universities concentrate on basic research, where free dissemination of knowledge is most important, and firms concentrate on R&D activities over which property rights should be strong. It would be very unwise for university researchers to perform proprietary research for private firms, yet many collaborations between business firms and universities or teaching hospitals are now taking precisely that direction. Universities in search of additional funding are increasingly seeking out arrangements under which they give up some of the traditions of open science and in effect become research subcontractors employed by private firms. Because the industry investment boards can solve, or at least mitigate, the free-rider problem, they can support universities and help set the research agenda without endangering the free exchange of ideas.

Closer interaction between firms confronted with practical problems and researchers pursuing fundamental questions may lead not just to bigger economic benefits but also to better basic science. For example, solid state physics was enormously stimulated by an attempt at Bell

Labs to solve a reliability problem with vacuum tube amplifiers. The labs had scientists on staff who knew some of the basic properties of semiconductors from solid state physics. They used this knowledge as a basis for their experiments leading up to discovery of the transistor. This discovery fed back into the basic research and training activities of solid state physicists and completely transformed the field. High temperature superconductors and opto-electronic devices are just two of the many underexploited opportunities that pose as many challenges for basic scientists as they do for practical problem solvers. These areas are at least as promising for the pursuit of basic knowledge as is the search in high energy physics for “the God particle,” using ever more expensive particle accelerators.

In summary, issuing the perennial call for a massive increase in federal support for existing basic research programs at universities is a totally unrealistic approach to a technology strategy. It is politically unrealistic and would not by itself generate much in the way of economic benefits. It might not even produce better basic science. The real challenge will be to keep government support for basic science from falling. At the same time, universities will feel growing pressure to contribute more directly to economic performance. Unless these pressures are channeled in a constructive way, they could undermine much of the strength of the basic research system, without achieving the desired practical benefits.

R&D Tax Credit

On efficiency grounds the optimal R&D tax credit would be so generous in its definition of R&D activity that no corporation would pay any income tax. Because of the widely acknowledged distortions that arise from the double taxation of capital income (once as corporate income, then as individual income), the efficiency gains from abolishing the corporate income tax would be large.

The United States has a corporate income tax for political reasons, not on grounds of economic efficiency. The potential for substantial increases in the generosity of the R&D tax credit is therefore limited by the lack of substitute funding sources and sensitivity to the political risks of giving “tax preferences to corporations instead of people.” If politicians continue to insist on raising revenue from the corporate

income tax, an argument can be made, in principle, for taxing R&D capital less heavily than physical capital because R&D has a higher social rate of return. In practice, the crucial issue is whether the tax authorities are able to distinguish R&D spending from other kinds of spending. Corporations will surely get better at redefining what they call R&D in response to bigger tax incentives for doing so. Over time, the revenue cost of each additional dollar of new R&D generated through this mechanism will grow. As a result, serious practical limits will be placed on the size of the subsidy that can be provided in this way as long as substantial revenue is collected through the corporate income tax.

Compared with other, more specific technology policies, the R&D tax credit is relatively safe politically, which no doubt explains why it has attracted broader political support than direct subsidies to firms. In effect, the tax credit specifies an aggregate level of subsidies for R&D, and then it solves the delegation problem by letting individual firms select their own R&D projects.

Whatever the decision on the constrained, second-best level of the corporate tax rate and the R&D tax credit, it should be clear that R&D spending by private firms (like government-funded pure basic scientific research) is a complement to the kind of industry-specific public goods that the industry investment boards would provide. Taken together, these three mechanisms (private R&D in firms, government support for basic research, and industry support for industry-specific public goods) would cover most of the spectrum of new ideas illustrated in the non-rival column in table 1.

Infrastructure

Increased funding for infrastructure as a technology policy has at least two serious strikes against it. First, no economic case can be made for infrastructure as a good that requires collective action. Very little infrastructure is a nonrival good. At best infrastructure is what is called a club good in public finance, and these goods can be efficiently provided by decentralized action in markets. Except for the bargaining problem over rights-of-way that is readily solved by the powers of eminent domain, no collective action is needed to provide infrastructure. (The state could use its powers of eminent domain to assist in land

acquisition where necessary but still rely on private sector decision processes in all other areas.) If anything, governments should be moving away from government-subsidized infrastructure and toward greater *reliance on market-like mechanisms such as congestion charges on roads and competition among providers of basic telephone service*. Localities or even private firms should be allowed to select and fund infrastructure projects.

The second problem is that federal funding for infrastructure is rife with the kind of waste inherent in a system that lacks any degree of delegation in project selection and funding. The government continues to provide support for “mass transit” that defies all economic logic. (My favorite example is the subsidy for the buses in Vail, Colorado, that take skiers through town and between ski lifts.)

The dollars at stake in decisions on popular infrastructure projects are truly frightening. A plausible estimate for the cost of roadbed construction for magnetic levitation trains is \$60 million a mile. Just putting in the track for a train that runs from Boston to Chicago could cost \$60 billion, or about thirty years worth of funding for the National Science Foundation.

Dual-Use Technology

Many advocates of a more aggressive technology policy have pointed to the apparent successes at DARPA as a model for a more activist role for the government. This approach should set off all kinds of warning bells about serious political risks. The most important problems pertain to funding, delegation, divided government, selection, inertia, and the robust reflection test.

The funding problem is self-explanatory. Budget constraints mean that the federal government is highly unlikely to make available enough funds to undertake a technology policy initiative that could make much difference on a national scale. The delegation problem has already been noted in the discussion of the contract for research on flat panel display screens. ARPA will be under increasing pressure to conform to congressional demands, especially if its budget grows. ARPA projects would then offer the best opportunity for every member of Congress to provide an identifiable local project. ARPA projects would soon look like infrastructure, earmarked pork-barrel science projects, and military pro-

curement. Asymptotically, every technology project would have part of its spending located in every congressional district, and the selection process for projects would become increasingly rule-bound and bureaucratic.

In its proposal on technology policy, the National Academy of Sciences Panel on the Government Role in Civilian Technology proposed that an independent civilian technology corporation be created with a one-time grant from the government of \$5 billion.²¹ This proposal reflects the panel's appreciation of the importance of the delegation and divided government problems, but it cannot be regarded as a true solution. A private corporation would in principle be free to make independent decisions and would not be subject to vetoes and second-guessing by multiple actors in the government. But the proposed \$5 billion initial grant is far too small to make a permanent contribution to a national technology policy. (This relatively low figure was presumably selected in the hope that Congress would fund it. Even at this low level, funding now seems extremely unlikely.) The interest income on \$5 billion would be about \$250 million a year. Compare this with the \$5 billion or more that IBM and GM each spend on R&D annually. Surely, the proponents of the Civilian Technology Corporation are assuming that if the corporation succeeds, it could go back for more funds. But that means the corporation would have to be highly responsive to the wishes of members of Congress from the start.

The comparison with IBM and GM also raises a caution about the selection problem and institutional inertia. Some proponents of the ARPA, or government planning, approach argue that industry leaders are not competent to make the right decisions about technology in the way a few key people at DARPA are alleged to have done. The countervailing argument is that putting all of the public's technology resources into a single hierarchical, bureaucratic structure leaves no protection if that organization stops producing valuable results. None of the existing proposals put any competitive pressure on ARPA or an independent technology corporation. If the entire auto and computer industries had depended on GM and IBM for all of their innovation, cars and computers would now be much less useful than they are. That

21. Panel on the Government Role in Civilian Technology (1992, p. 94).

is true even though the research staffs in both corporations are highly talented and even though some discoveries of real value have emerged from both institutions (for example, catalytic converters from GM and high temperature superconductors at IBM.) There is no substitute for competitive pressure on any such organization.

Even if all of the other objections could be met, the technology agency approach is probably not the kind of policy that passes the reflection or robust reflection tests. It is all too easy to imagine that an ARPA or a well-funded civilian technology program could be used to subsidize something like the European Airbus program, which Americans now find so objectionable. Mechanisms for direct government subsidies to firms may not be something that the United States wants to encourage in other countries.

Research Joint Ventures

Interest in the United States in research joint ventures seems to have been motivated primarily by the success of the VLSI (very large scale integration) joint venture among semiconductor firms in Japan, an example that is now understood to have been unusually successful and not representative of the typical outcome. Research joint ventures suffer from the obvious economic problem that they cannot solve the free-rider problem and achieve free dissemination of the results at the same time. What has been attempted in the United States, for example at SEMATECH, is a mixture of some degree of insider advantage that is promised to participants in the joint venture together with federal matching funding to sweeten the pot. One half of SEMATECH funding was provided by DARPA, so if this model were to be adopted more widely, all of the concerns about delegation and divided government that apply to DARPA would extend to the government-assisted joint venture.

But if government funds are not used, in any instance where several firms in an industry find it in their interest to work together on a research project, they presumably will be motivated to do so by the availability of important special advantages that accrue only to members of the joint venture. Such cases should raise serious concern that collusion in research will lead to product market collusion. A research joint venture

could, for example, become the vehicle for preventing new entry into an industry or for reducing total industry expenditures on research.²² Compared with the proposal outlined above, in which all joint effort has to be directed at activities that are freely available to any current or future participant in the industry, the research joint venture seems to hold out much greater risk of reduced total research and product market competition than do industry investment boards.

Managed Trade: Import Restrictions

Managed trade comes in two flavors, import restrictions and market access requirements. Import restrictions have the appeal that they are, in effect, self-financing. Import quotas, for example, can generate additional revenue for protected firms without imposing an explicit tax. Import restrictions suffer from almost every other conceivable problem. Some firms will support this option out of naked self-interest, but virtually no economist is willing to try to make a serious intellectual case for these measures.

Managed Trade: Market Access Requirements

Proposals that the U.S. government demand adequate access to foreign markets are more problematic. The purely economic case here is stronger than many economists who support free trade would like to admit. Market access is a necessary condition for effective protection of intellectual property rights. Suppose, for example, that a foreign country enforces a ten-year quarantine on all imported foreign software to check for software viruses. (Sounds plausible, right?) Domestically produced versions of word-processing software, spreadsheet programs, and other kinds of applications that are close copies of the versions available on world markets would no doubt be written and sold in this country. The effect for U.S. software firms would be as if the foreign firms simply sold bootleg versions of their programs. Formal copyright protection is worth nothing if a firm is deprived of the right to sell its goods.

It is not clear that market access restrictions would be in the self-interest of the foreign country. It depends on how costly the copying

22. For a discussion of these issues, see Katz and Ordover (1990).

process is and on what effects the protection has on domestic incentives. The experiences in Brazil, with its failed computer industry, and in Japan, with its (allegedly) protected personal computer industry, do not speak well for the long-run success of this strategy. If a protected domestic market were truly the key to export success, NEC and the Brazilians should have taken over the market for personal computers in the United States, but they have not. For internal political reasons, many countries may nevertheless restrict market access. The United States must therefore decide how it should respond to these restrictions.

Some advocate unilateral punitive measures, but this approach fails the robust reflection test. If market access requirements were given the same status as intellectual property rights and were monitored and enforced through a multilateral body, they would be beneficial. But a unilateral decision by the United States to adopt punitive measures as it sees fit could ultimately have very serious negative effects on the institutions that now support free trade. This would effectively condone the principle that countries can take unilateral action when they are unhappy with trade outcomes. Many other countries would soon start instituting punitive measures that are supposedly intended to punish market restrictions elsewhere but that are, in fact, pure income transfers to powerful domestic groups.

Moving ahead with unilateral punishments would be like announcing an intent to drive through red lights whenever no cars are coming on the cross street. We may believe that no harm will come from our actions if we use our new powers judiciously, and we may even be willing to let others do what we do if they promise to be as responsible as we are. The ultimate outcome is nevertheless likely to be one that we regret.

A useful comparison would be to contemplate the consequences of repealing the interstate commerce clause in the U.S. Constitution and letting individual states take punitive action whenever another state restricts market access. (Many state governments in the United States currently do try to give the same preference to local firms that the Japanese government apparently gives in procurement.)

Economists would almost universally agree that giving each state the power to interfere with interstate commerce would not be the way to deal with market access problems within the United States. The same argument presumably suggests that the United States should work to

strengthen the General Agreement on Tariffs and Trade and the available multilateral institutions in the hope of moving trade among countries closer to the substantial, but not quite perfect, freedom that now prevails among the states of the United States.

Nevertheless, attacks on restricted market access will continue to be a very attractive political strategy for the federal government. Such attacks can be used to aid a small number of firms with high visibility. They do not require budget expenditures, they typically raise little domestic political opposition, and a confrontation with bad guys from another country always plays well in politics. (The war over the Falkland Islands may indeed have been like two bald men fighting over a comb, but it did wonders for the popularity of the Thatcher government.) Of all the policies described here, saber rattling and punitive measures designed to open foreign markets are the most likely candidates for a technology strategy in coming years.

Conclusions

In one of the last things he wrote, the late George Stigler concluded an essay on monopoly with these words: “The merits of *laissez-faire* rest less upon its famous theoretical foundations than upon its advantages over the actual performance of rival forms of economic organization.”²³ Any discussion of technology policy should take these words as both a warning and a challenge. The warning is that even after admitting all of the deficiencies of economic markets, one must acknowledge what experience has so clearly demonstrated—that most of the familiar political alternatives are far less efficient mechanisms for allocating resources than the market is. The challenge for economists is to understand why markets perform well and then to build upon their strengths. The problem with the classical description of *laissez-faire* is its suggestion that the best of all possible arrangements for economic affairs has already been discovered and that it requires no collective action. The lesson from economic growth is that collective action is very important and that everything, including institutions, can always be improved.

23. Stigler (1993, p. 402).

The reality of economic life is that the pure individualistic market exchange invoked in the usual defense of *laissez-faire* has been supplanted by many institutions that allow collective action and improve on pure individualistic market exchange. Perhaps the most important such institution is the limited liability corporation, an invention that would not have been viable without the emergence of supporting regulatory and legal institutions. Other examples of institutional innovation include private universities, copyrights and patents, corporate research laboratories, and peer-reviewed competitive grants for research.

Each of these institutions represents an attempt to take advantage of the opportunities for mutually beneficial coordinated action that are presented by the world in which we live. Each is consistent with the operation of markets and supports success in the market as the final test of economic activity. Each builds on or emulates at least some of the strengths of the market: the push to do better that comes from competing against rivals, the pull that comes from opportunities for individual gain, the diversity that comes from having many different individuals and organizations working in parallel to achieve a given end, and the discipline that comes from clearly enforced criteria for success and failure. From this broader perspective, the “theoretical foundations” to which Stigler attached relatively little importance—the notions of price taking and tangency conditions—are only one small part of the picture.

This paper has tried to place the discussion of technology policy in this larger context. It has tried to shift the discussion away from narrow questions about whether the market or the government is better. It has avoided conjectures about which sectors should be given government support or about what are the most important “critical technologies.” Instead, it has suggested that the debate focus on the underlying processes that lead to effective institutional arrangements.

To a large extent, market competition has become a process for selecting ever better corporate institutions that can channel the energies of large numbers of people toward the production of ordinary private goods. This paper describes a parallel process for selecting ever better institutions that can channel the energies of large numbers of people toward the production of the new discoveries that drive economic growth.

The process outlined above combines government-like aspects (man-

datory taxes that solve free-rider problems) with market-like mechanisms (free entry, competition among different institutions, and decisions that are ultimately grounded in opportunities for profit). It is designed to fill the gap between private and public support for R&D. The essence of the proposal is to empower the firms in different industries, giving them the tools needed to solve the collective action problems inherent in providing industry-specific public goods. The proposal protects product market competition among firms and encourages competition among the organizations that would provide industry-specific public goods such as research or training. The argument is premised on the idea that long-run growth depends on our ability to discover innovative ways to arrange the limited stocks of objects. It concludes that the prospects for growth could be enhanced by searching for equally innovative ways to arrange our institutions.

Comments and Discussion

Comment by Zvi Griliches: This is an exciting paper. It is nice to see somebody young tackling big policy issues. The paper builds on an analysis of the past, but it is basically about the future. It is suggesting new mechanisms and new institutions to solve pressing economic problems. It takes as its historical example the “invention” of land grant colleges and the development of Bell Labs. It sees unlimited opportunities for investment in both basic and applied research and is looking for mechanisms that could provide the financing that would launch the next expeditions into the wide unknown.

The two basic assumptions of the paper are, first, that the research opportunities are out there and, second, that without new financing institutions, we may not get there. I believe that Romer is right on both counts but that the evidence for the first is not as overwhelming as he assumes it to be and that the particular institutional innovation may not work as well as advertised—but nothing ventured, nothing gained.

I will focus on the second point first. I see two difficulties with Romer’s proposal. The first deals with the incidence of benefits from and the costs of technical change in an industry. Implicitly, a high elasticity of demand is assumed for the products of the various industries. Because the benefits from this research will be public, they are likely to be competed away and not redound to the benefit of the original investors in these boards. Moreover, if demand were inelastic, as in agriculture, such technical change could reduce the rents to the scarce resources in these industries and might be counterproductive from the private point of view. Thus, industries might not vote for such a levy if they see through it.

Second, if there are several boards and I am a small producer, I may choose to contribute my share to a peculiar board, specializing, for example, in supporting economic research at the University of Chicago, on the assumption that the big guys will give to the important research purposes anyway and that my contribution to those boards is unlikely to affect the industry's fate by much. In other words, allowing for alternative boards and for freedom of choice among them may still leave free-rider problems, even within this framework. Nevertheless, the proposal may be worth trying. People may not be as selfish as economists say they should be.

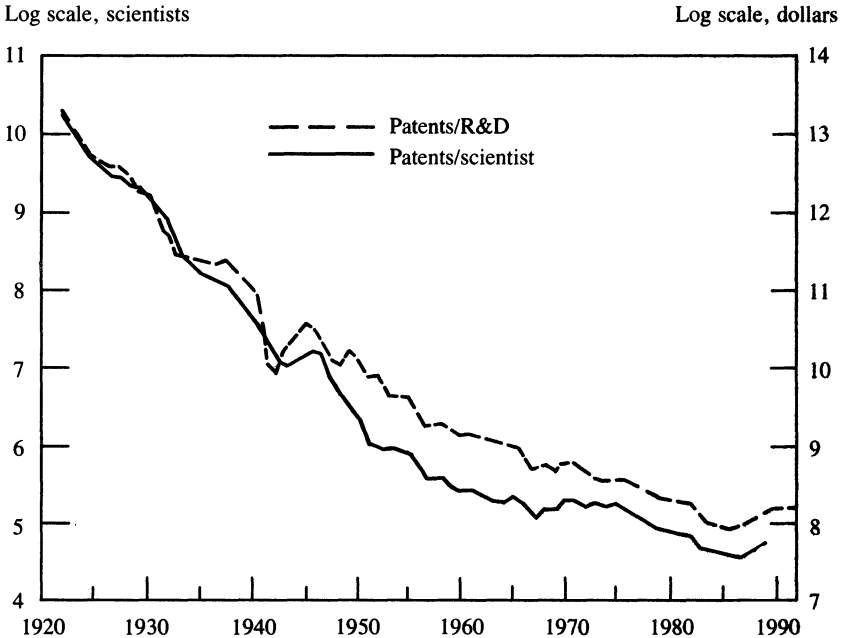
The more difficult notion is the one of the "endless frontier," the absence of diminishing returns to research. Because I have been worrying about this lately, I will focus the rest of my remarks on what we do not know about it.

One can ask two related questions about this. Are our recent "slow-growth" problems caused by a shortfall of research and development? And will encouraging more R&D be helpful here? I don't think that the sharp recession of the mid-1970s and the subsequent slow growth were the result either of the decline in the rate of aggregate R&D investment that occurred in the late 1960s and early 1970s or of a fall in the "potency" of these expenditures. The timing and magnitudes are all wrong.¹ It is probably a good thing to pursue more R&D in any case, for reasons of long-term growth, but not as a solution to our current problems. (Of course, Romer is not saying that his proposal will solve these problems, but it is the current economic situation that gives his proposal its resonance.) We are not currently constrained by our production possibilities frontier. Much existing knowledge is not being used yet. Much more diffusion of computers and learning about them is still in store. And the economy could be made much more efficient by eliminating a variety of barriers to its functioning. Nevertheless, R&D is one of the relevant margins, and we should be exploring it.

But before we encourage more R&D by creating new institutions, we have to face the crude evidence of diminishing returns to it. Figure 1 plots the ratio of patents to company-financed R&D for the last seventy years. If patents are taken as an indicator of inventive output, then the news is not all that good. Bronwyn Hall's paper in this volume

1. See Griliches (1988) for an elaboration of this argument.

Figure 1. Domestic Patent Applications per Company-Financed R&D in Industry (1972 dollars) and per Scientist and Engineer, 1921–92.



Source: Bureau of the Census, *Historical Statistics of the United States, Colonial Times to 1970*, pt. 2 (Department of Commerce, 1975), chap. W; Patent and Trademark Office, *Technology Assessment and Forecasts* (Department of Commerce, March 1977), app. A; National Science Board, *Science and Engineering Indicators*, 1991, NSB 91-1 (Washington); and Patent and Trademark Office releases. Number of applications by residence of inventor for 1940–59, unpublished memorandum of P. F. Fredrico, Patent and Trademark Office, January 18, 1961; for 1960, *Journal of the Patent Office Society*, 44 (February 1964), p. 168; for 1961–62, Commissioner of Patents and Trademarks, *Annual Report, 1966*, p. 26; for 1963–91, Patent and Trademark Office releases; extrapolated back, before 1940, by the number of total applications. Data on company R&D and the number of scientists and engineers in industry from the same sources, updated from NSF Reports 92-307, 92-330, and 92-331. Deflated by R&D price index from Griliches, "Comment on Mansfield," in *R&D, Patents, and Productivity*, (University of Chicago Press, 1984) p. 149, and Bronwyn H. Hall and others, "The R&D Master File Documentation," Technical Paper 72 (NBER, December 1988), updated.

Note: Left scale is the log of the number of domestic patent applications per 1,000 R&D scientists and engineers. Right scale is the log of the number of domestic patent applications per \$1 million of company-financed R&D (in 1972 dollars).

shows a significant decline in the market’s recent valuation of R&D. Evidence claiming to show exhaustion of inventive opportunity was presented by Baily and Chakrabarti at these same meetings some years ago, and Evenson looks at the international evidence on patenting and comes to similar conclusions.²

I myself do not think these facts need to be interpreted so pessimistically, but before we urge significant additional investments in R&D,

2. Baily and Chakrabarti (1988); Evenson (1993).

we need to convince ourselves that this would be a good investment. My own optimistic reading of these same data is based on scattered evidence that the economic “meaning” of patents has changed. This is the only reasonable way to interpret the fact that ratios of patents to R&D have been falling consistently, until quite recently, through good times and bad. If we did not worry about it in the 1950s and 1960s, when total factor productivity growth was rapid, why should we worry now? Moreover, even while—according to Hall’s estimates—the market’s valuation of new R&D was falling sharply, firms were expanding their R&D investments at rates not seen since the 1950s! Did companies know something that the market did not, or were they all lemmings? At the moment we face two contradictory facts: domestic patent applications started growing sharply in the late 1980s, indicating some revival of inventive activity, while the growth in companies’ real R&D expenditures slowed to a crawl and may have actually turned negative. But overall, if the numbers are to be believed, we are not investing less in R&D today, relative to the size of the economy, than we did in the peak years of the 1960s. It is a puzzlement.

An argument can be made that precisely because we may have been facing the exhaustion of inventive opportunities, we need more R&D, but R&D of a special kind, the basic kind, the kind that would “re-charge” the pool of knowledge and increase the effectiveness of the rest of it, which is spent on “D,” rather than “R.” We need additional investment in science and basic research to make the run-of-the-mill company’s R&D more productive. In this sense, Romer’s new institutions, specializing in basic and “generic” industrial R&D may be just what the doctor ordered. But what we want is basic research and basic results that will not be appropriated by individual firms, and I am not sure that we will succeed in convincing such firms to tax themselves for the public good.

I do want to dissent from Romer’s view that the fact the molecules can be arranged in an infinity of ways implies an “endless” frontier for R&D exploration. It may not be feasible to discover the work of a new Shakespeare in the random typing of monkeys, even if it is there. More seriously, the notion that there may be no diminishing returns to research in the long run, and the associated notion of a permanent exponential growth of real income per capita, is a dream full of hubris, the notion that man is God-like and not subject to serious constraints

or hard choices. Many years ago, in commenting on the “old” growth theory, Arrow remarked that “eternal exponential technological growth is just as unreasonable as eternal exponential population growth.”³

The fact that the frontier may not be endless need not be taken pessimistically. First, we are still very far from being there. And, second, much is left to do to improve the situation as it is, including the diffusion of already-known technologies and the dismantling of the many manmade barriers to economic efficiency. Technological change is still endogenous, even if it is not endless, and although the outcomes of science and research are largely out of our hands, better economic understanding still has much to contribute. Schmookler, in one of the last things he wrote, complained about the state of this field: “. . . it may leave scientific and technological progress unexplained, which is unsatisfactory; alternatively, it explains scientific and technological progress as self-generated, which is dreadfully wrong. Such a view is wrong both scientifically and morally. It is wrong scientifically because . . . it is simply not true. It is wrong morally . . . because it deprives man generally of any sense of responsibility for the course of social and economic development. For if [technological change is endogenous], then all men must accept some measure of responsibility for what happens next.”⁴ It is admirable that Romer is willing to do just that.

General Discussion: In evaluating the potential of the author’s plan, several participants raised examples of past experience with cooperative research efforts. Pointing out that few firms have formed research joint ventures in the decade since it first became permissible under antitrust laws, Timothy Bresnahan wondered if there really is an unfulfilled demand for more cooperative research. He also noted that existing cooperative institutions in the form of standards committees (such as ANSI) are abandoned by many in information technology industries when these committees attempt to coordinate the direction of technological progress through anticipatory standard setting. Bresnahan considered the rules imposed by standards committees to be quite modest and argued that the author’s seemingly more restrictive institutional arrangements could prove even less popular.

3. Arrow (1969, p. 34).

4. Schmookler (1972, p. 84).

In contrast, Albert Link said that experience with SEMATECH, the cooperative research venture of the semiconductor industry, shows that cooperation can yield important benefits. Despite others' criticisms of that organization, Link said, it has ensured that American companies are competitive in their production of semiconductor manufacturing equipment, such as wafer polishing tools.

Link, however, questioned whether the author's proposal would succeed in increasing basic scientific research. He maintained that existing cooperative arrangements with universities are becoming more commercially oriented and now almost mirror research that is performed at private firms. Link nonetheless praised the proposal's concern with providing a forum for the more efficient evolution of standards, saying that this could lead to a more rapid diffusion of technology throughout a given industry.

Linda Cohen argued that the coercive part of the author's proposal might make it more viable than existing forms of noncoercive cooperative arrangements. She maintained that noncoercive cooperative agreements work well only in regulated and monopolistic industries (and, for unique reasons, in agriculture), and that they tend to break down under competitive conditions. She noted that when the electric industry became more competitive, Southern California Edison, one of the largest and most innovative utilities, dropped out of the Electric Power Research Institute because the utility no longer wanted to share technology. Using SEMATECH as an example, she argued that when noncoercive cooperation involves the public sector, the results of the research are forced to be more public, allowing many firms to opt out of the arrangement and get a free ride. In competitive markets, she concluded, coercion forcing 100 percent of participation may be the only feasible way to bring about sustained cooperative research.

Michael Katz argued that the government role in the author's proposed arrangement would end up being far more intrusive than the author assumes. He said that government would become integral in determining which firms fall within an industry and which do not. He also suggested that the author's proposal to make imported goods subject to the mandatory industry taxes could lead to abuse if other countries adopted this proposal as well. If his plan were effected in a country such as France, Katz said, nationalized French firms might be able to avoid paying their taxes through offsetting government subsidies, plac-

ing American and other foreign firms subject to the tax under an unfair burden.

Several participants discussed the author's presumption that almost limitless potential gains can be had from pushing out the knowledge frontier. Robert Gordon argued that diminishing returns present a serious constraint to the benefits of technical progress. He suggested that such diminishing returns can be observed by looking at the effects of the technological changes of the past century. Such changes have been transforming daily life over this period, but the rate of transformation has been slowing down. Martin Baily, in contrast, argued that the significant externalities connected with additional investments in research and development make a strong case for pushing out the knowledge frontier.

Robert Hall disagreed with the author's contention that providing extremely strong property rights over nonrival goods would lead to inefficiencies and lost opportunities for additional technological progress. Using the computer industry as an example, he noted that IBM has been liberally granted licenses for the use of its key patents. Bresnahan concurred with Hall's assessment and added that for the commercialization of information technology, the most important mechanism for coordination and reuse is the appointment of a monopoly vendor.

References

- Arrow, Kenneth J. 1969. "Classificatory Notes on the Production and Transmission of Technological Knowledge." *American Economic Review* 59 (May): 29–35.
- Baily, Martin Neil, and Alok K. Chakrabarti. 1988. *Innovation and the Productivity Crisis*. Washington, D.C.: Brookings.
- Bush, Vannevar. 1945. *Science: The Endless Frontier: A Report to the President on a Program for Postwar Scientific Research*. Washington, D.C.: National Science Foundation.
- Carey, John. 1993. "Industrial Policy or Industrial Folly?" *Business Week* (May 17): 38.
- Coase, R. H. 1988. *The Firm, the Market, and the Law*. Chicago: University of Chicago Press.
- Coe, David T., and Elhanan Helpman. 1993. "International R&D Spillovers." Working Paper 4444. Cambridge, Mass.: National Bureau of Economic Research. August.
- Cohen, Linda R., and Roger G. Noll. 1991. *The Technology Pork Barrel*. Washington, D.C.: Brookings.
- Crease, Robert P. 1991. "Yesterday's Pork Projects: Where Are They Now?" *Science* 254 (November 1): 641–43.
- Ergas, Henry. 1987. "Does Technology Policy Matter?" In Guile, Bruce R., and Harvey Brooks, eds. *Technology and Global Industry: Companies and Nations in the World Economy*. Washington, D.C.: National Academy Press.
- Evenson, R. E. 1993. "Patents, R&D, and Invention Potential: International Evidence." *American Economic Review* 83 (May): 463–68.
- Griliches, Zvi. 1958. "Research Costs and Social Returns: Hybrid Corn and Related Innovations." *Journal of Political Economy* 66 (October): 419–31.
- . 1988. "Productivity Puzzles and R&D: Another Nonexplanation." *Journal of Economic Perspectives* 2 (Fall): 9–21.
- . 1992. "The Search for R&D Spillovers." *Scandinavian Journal of Economics* 94 (Supplement): S29–S47.
- Katz, Michael L., and Janusz A. Ordover. 1990. "R&D Cooperation and Competition." *Brookings Papers on Economic Activity, Microeconomics: 1990*: 137–203.
- Lichtenberg, Frank R. 1992. "R&D Investment and International Productivity Differences." Working Paper 4161. Cambridge, Mass.: National Bureau of Economic Research. September.
- Mansfield, Edwin, J. Rapoport, A. Romeo, S. Wagner, and G. Beardsley. 1977. *The Production and Application of New Industrial Technology*. New York: Norton.
- Merges, Robert P., and Richard R. Nelson. 1992. "On Limiting or Encour-

- aging Rivalry in Technical Progress: The Effect of Patent Scope Decisions.” University of California (Berkeley) Center for Research and Management.
- Nelson, Richard R. 1983. “Government Support of Technical Progress: Lessons from History.” *Journal of Policy Analysis and Management* 2 (Summer): 499–514.
- Office of Technology Assessment, U.S. Congress. 1991. *Federally Funded Research: Decisions for a Decade*. Washington, D.C.: GPO.
- Panel on the Government Role in Civilian Technology, Committee on Science, Engineering, and Public Policy, National Academy of Sciences. 1992. *The Government Role in Civilian Technology: Building a New Alliance*. Washington, D.C.: National Academy Press.
- Powers, Nicholas J. 1990. “Federal Marketing Orders for Fruits, Vegetables, Nuts, and Specialty Crops.” Agricultural Economic Report 629. U.S. Department of Agriculture, Economic Research Center, Washington, D.C. March.
- Rosenberg, Nathan, and Richard R. Nelson. 1993. “American Universities and Technical Advance in Industry.” CEPR Publication 342. Center for Economic Policy Research, Stanford University. March.
- Savage, James D. 1992. “Trends in the Distribution of Apparent Academic Earmarks in the Federal Government’s FY 1980–92 Appropriations Bills.” *CRS Report for Congress*. 92-726STR. GPO.
- Schmookler, Jacob. 1972. *Patents, Invention, and Economic Change: Data and Selected Essays*. Cambridge, Mass.: Harvard University Press.
- Schultz, Theodore W. 1953. *Economic Organization of Agriculture*. New York: McGraw-Hill.
- Stigler, George. 1993. “Monopoly.” In Henderson, David R., ed. *The Fortune Encyclopedia of Economics*. New York: Warner Books.
- Williamson, Oliver E. 1975. *Markets and Hierarchies: Analysis and Antitrust Implications*. New York: Free Press.
- Wilson, James Q. 1989. *Bureaucracy: What Government Agencies Do and Why They Do It*. New York: Basic Books.