This PDF is a selection from an out-of-print volume from the National Bureau of Economic Research

Volume Title: Innovation Policy and the Economy, Volume 1

Volume Author/Editor: Adam B. Jaffe, Josh Lerner and Scott Stern, editors

Volume Publisher: MIT Press

Volume ISBN: 0-262-60041-2

Volume URL: http://www.nber.org/books/jaff01-1

Publication Date: January 2001

Chapter Title: Should the Government Subsidize Supply or Demand in the Market for Scientists and Engineers?

Chapter Author: Paul M. Romer

Chapter URL: http://www.nber.org/chapters/c10781

Chapter pages in book: (p. 221 - 252)

7

Should the Government Subsidize Supply or Demand in the Market for Scientists and Engineers?

Paul M. Romer, *Graduate School of Business, Stanford University*

Executive Summary

This paper suggests that innovation policy in the United States has erred by subsidizing the private sector demand for scientists and engineers without asking whether the educational system provides the supply response necessary for these subsidies to work. It suggests that the existing institutional arrangements in higher education limit this supply response. To illustrate the path not taken, the paper considers specific programs that could increase the numbers of scientists and engineers available to the private sector.

Preface

My son attends an undergraduate institution that specializes in science and engineering. A degree from this school will cost more than $100,000 in tuition and 4 years of his time. For parents and students who contemplate an investment of this magnitude, the school provides information about labor market outcomes for its graduates. On its web site, the school provides the median salaries for students who accepted jobs after graduation and the Ph.D. completion rates for students who go on to graduate school. If you search, you can see the entire distribution of outcomes—a listing for each student of the starting salary or graduate school in which they enrolled.

If my son pursues a doctoral degree after he graduates, he will have to make an even larger investment. Net of the various sources of support that are available to graduate students, the direct tuition cost of a doctorate will probably be less than the cost of his undergraduate degree. He will, however, have to invest another 4 to 8 years of foregone earnings, which will be substantially higher once he completes his undergraduate degree.

His college is unusual. Most undergraduate institutions do not provide any useful information about labor market outcomes for degree recipients. Yet as

this example shows, it is perfectly feasible for a school to provide this kind of information. Given the stakes, it is even more surprising that many graduate programs in science and engineering also fail to provide this kind of information to prospective students. The paucity of information is obvious to anyone who is familiar with the graduate school application process, but to demonstrate it more formally, I asked a research assistant to begin the application process for the top 10 graduate departments of mathematics, physics, chemistry, biology, computer science, and electrical engineering in the United States. (The rankings were taken from US News and World Report.) For comparison, I also asked him to begin the application process to the top 10 business and law schools.

In response to his 60 initial requests for information from the science and engineering programs, he received not one response giving information about the distribution of salaries for graduates, either in the initial information packet or in response to a follow-up inquiry from him. In contrast, he received salary information for 7 of the 10 business schools in the application packet, and in response to his second request, he was directed to a web page with salary information by one of the three nonrespondents from the first round. (It is possible that the information could have been found on the web page for the other two business schools, but to maintain consistency in the treatment of the different programs, he did not look for more information if a school did not respond with directions about where to get it.) Four out of the 10 law schools gave salary information in the application packet and three more of them directed him to this information in response to a second request.

I. Introduction

The most important economic policy question facing the advanced countries of the world is how to increase the trend rate of growth of output per capita. In the middle of the 20th century one might have argued that preventing depressions was the more urgent challenge, but at least in the advanced countries of the world, progress in macroeconomic stabilization policy has reduced the threat of a paralyzing economic collapse and even reduced the frequency of mild recessions. In this environment, the lure of better growth policy is compelling. If an economy can increase its trend rate of growth by even a small amount, the cumulative effect on standards of living is too big to ignore.

Many scholars and policymakers are convinced that during the 20th century, rapid technological progress in the United States drove the un-

precedented growth in output and standards of living we enjoyed. In addition, they believe that this rapid rate of technological change was fostered by a publicly supported system of education that provided the essential input into the process of discovery and innovation—a steady flow of people trained in the scientific method and in the state of the art in their area of specialization.

If this interpretation of our recent past is correct, it follows that any proposal for achieving an even higher trend rate of growth in the United States should take full account of the detailed structure of our current system of higher education for natural scientists and engineers. Policymakers must recognize that these institutions exhibit puzzling features such as those described in the preface—an almost total lack of information on future market opportunities for students who enter their programs.

Unfortunately, in the last 20 years, innovation policy in the United States has almost entirely ignored the structure of our institutions of higher education. As a result, government programs that were intended to speed up the rate of technological progress may in fact have had little positive effect. We have undertaken major spending initiatives in the area of innovation policy, our most important area of economic policy, without subjecting their economic assumptions to even a cursory check for logical coherence or factual accuracy.

In what follows, I will point to the fundamental conceptual flaw behind the government programs that have been used in the last 20 years to encourage innovation in the private sector. These programs try to stimulate the demand for scientists and engineers in the private sector. To succeed, they depend on a positive supply response that the educational system seems incapable of providing. I will also describe a class of alternative policy programs that could be used to fill the gap created by an exclusive reliance on demand-side subsidies. These alternative programs return to an early style of government policy, one that works directly to increase the supply of scientific and engineering talent.

Section II below starts with a quick recapitulation of the reasons why decisions concerning innovation policy are so important for the economic well-being of future generations. Section III shows how a demand-side innovation policy such as a tax credit for research and development or a program of research grants affects the market for scientists and engineers. It shows why even a well-designed and extremely generous program of this kind will fail to induce more

innovation and faster growth if the educational system does increase supply in response to changes in wages. Section IV provides an overview of trends in the supply of scientists and engineers. Sections V and VI look at the market for undergraduates and for Ph.D. recipients respectively. Section VII summarizes and interprets the evidence from these sections.

One of the surprising features of the political debate surrounding demand-subsidy policies is its narrow focus. Few participants in this debate seem to have considered the broad range of alternative programs. Section VIII tries to broaden the debate by suggesting feasible policies that could be considered. More specifically, it outlines a general process that policymakers could adopt for formulating growth policy. This process starts by distinguishing between goals and programs. To be specific, this section outlines four general goals that could guide the formulation of growth policy. The first possible goal that policymakers might adopt would be to target a specific increase in the number of students who receive undergraduate degrees in the natural sciences and engineering. A second would be to encourage more innovation in the graduate training programs that our universities offer to students who are interested in careers in science and engineering. A third would be to preserve the strength of our existing system of Ph.D. education. A fourth would balance amounts that the federal government spends on subsidies for supply and demand of scientists and engineers.

If policymakers in an economy were to adopt goals such as these, the next step would be to design specific programs that are intended to achieve these goals. In broadening the debate, this section also outlines three illustrative programs that policymakers could adopt to achieve these goals. The first is the introduction of training grants to universities that could be used to increase the fraction of undergraduates who receive degrees in natural science and engineering. The second is a system of exams that give objective measures of undergraduate achievement in natural science and engineering. The third is a new type of fellowship, backed by a substantial increase in funding, for students who continue their studies in graduate school.

The advantage of a process that separates goals from programs is that it establishes a natural way to evaluate any specific programs such as these. If the goals are precise and progress toward them can be quantified, then it should be easy to verify if any given program moves the economy closer to the goals. This makes it possible to experiment

GDP per Capita

Figure 7.1
Income per capita from 1870 to 1992 for the U.S., Britain, and Sweden.
Data are from Angus Maddison (1995).

with a variety of programs, to expand the ones that work, and to shut down the ones that do not.

II. The Importance of Technology Policy

A quick look at the data in figure 7.1 suggests that there must be policy choices, intentional or unintentional, that affect the trend rate of growth. Using data assembled by Angus Maddison (1995), this figure plots income per capita from 1870 to 1994 for the United States, Britain, and Sweden. Over the century-and-a-quarter of data presented there, income per capita grew in the United States at the rate of 1.8% per year. In Britain, it grew at 1.3% per year. In the beginning of the sample period, the United States was a technological laggard, so part of its more rapid growth could have come from a process of technological catch-up with Britain, which was at that time the worldwide technology leader. But even at the beginning of this period, it was clear that the United States was also capable of generating independent technological advances—for example, in the area of manufacturing based on the assembly of interchangeable parts. (See Rosenberg 1969 for an account of the reaction that the "American system of manufactures" caused in Britain by the middle of the 19th century.) Moreover, as the United States surged ahead of Britain in the 20th century, it maintained the

GDP per Capita

Figure 7.2
Same data as figure 7.1, on a logarithmic scale.
Data are from Angus Maddison (1995).

faster rate of growth that was apparent from the beginning. This is most clearly evident in figure 7.2. Because the data are plotted on a logarithmic or ratio scale, straight lines in the figure correspond to constant rates of growth. In the second half of the century, the rate of growth in Britain accelerated moderately. The rate of growth that had been initiated in the U.S. remained essentially unchanged. The policies and institutions in the United States made possible a trend rate of growth of income per person that was significantly faster than the trend that had pushed Britain into the position of worldwide leadership in the 19th century. Given the limited state of our knowledge of the process of technological change, we have no way to estimate what the upper bound on the feasible rate of growth for an economy might be. If economists had tried to make a judgment at the end of the 19th century, they would have been correct to argue that there was no historical precedent that could justify the possibility of an increase in the trend rate of growth of income per capita to 1.8% per year. Yet this increase is what we achieved in the 20th century.

The experience in Sweden suggests that even higher sustained rates of growth of income per capita can sometimes be possible. During the 50 years from 1920 to 1970, income per capita in Sweden grew at the much higher rate of 3% per year. Once again, this faster rate of growth could be due, at least in part, to the process of technological catch-up.

Moreover, growth in Sweden has slowed down considerably since 1970. Nevertheless, the experience in Sweden should at least force us to consider the possibility that if we arranged our institutions optimally, growth in the United States could take place at an even higher rate than that to which we have become accustomed. In the coming century, perhaps it will be possible to increase the rate of growth of income per capita by an additional 0.5% per year, to 2.3% per year.

The implications of a change of this magnitude would be staggering. For example, according to recent CBO estimates that were based on continuation of the historical trend rate of growth, in the year 2050 the three primary government entitlement programs—Social Security, Medicare, and Medicaid—will require an increase in government spending equal to about 9% of projected GDP. If the rate of growth over the next 50 years were to increase by 0.5% per year, GDP in 2050 would be 28 percentage points larger. By itself, faster growth could resolve all of the budget difficulties associated with the aging of the baby boom generation, and still leave ample resources for dealing with any number of other pressing social problems. And of course, the longer a higher growth rate can be sustained, the larger the effect it will have. By the year 2100, the additional 0.5% per year would translate into a GDP that is 1.65 times as large as it would otherwise have been.

Other types of evidence suggest that an increase in the rate of growth of 0.5% per year is not beyond the realm of possibility. In his survey of returns to investment in R&D, Zvi Griliches (1992) reports a wide range of estimates for the social return, with values that cluster in the range of 20% to 60%. Take 25% as a conservative estimate of the social return on additional investment in R&D. If we were to increase spending on R&D by 2% of GDP (and maintain the same rate of return on our investments in R&D—more on this in the next section) the rate of growth of output would increase by the hoped-for 0.5% per year. If the true social return is higher, say 50%, the extra investment in R&D needed to achieve this result would be correspondingly lower, just one additional percent of GDP. These estimates are also consistent with other estimates, which suggest that the level of resources currently devoted to research and development may be far below the efficient level. For example, after they calibrate a formal growth model to the results from micro level studies of the productivity of research and development, Chad Jones and John Williams (1998) calculate that the optimal

quantity of resources to devote to research and development could be four times greater than the current level.

There is another way to look at estimates of this kind, one that is closer to the spirit of the analysis that follows below. GDP in the United States is about $10 trillion dollars. One percent of this would be $100 billion per year in additional spending on R&D. If it costs $200,000 per year to hire and equip the average worker in this sector, this means that we would need to increase the stock of workers employed in R&D by roughly 500,000. The question that policymakers must confront if they are serious about increasing the amount of R&D that is performed is where these additional high-skilled workers will come from.

There is no certainty that growth would necessarily speed up even if we did undertake all the right steps in an effort to do so. There is ambiguity in the historical record, and even if there were not, there is no guarantee that relationships that held in the past will continue to hold in the future. Moreover, even in the best case, we should recognize that there might be substantial lags between the initiation of better policy and the realization of faster output growth. For example, one highly successful example of a government policy that did increase the rate of technological change was the creation of the new academic discipline of computer science in the 1960s. (See Langlois and Mowery 1996 for a discussion of the episode.) Even now, with the passage of 40 years, our sense of the magnitude of the payoff from this investment is still growing.

Notwithstanding all these caveats, a possibility, even a remote possibility, of a change as profound as another permanent 0.5% increase in the trend rate of growth in the world's leading economy ought to excite the imagination. Compared to this, even landing a man on the moon would seem a minor achievement. One would think that this kind of possibility would inspire us to try new things, to make every effort to understand what will work and what will not as we strive for this goal. By this kind of standard, the efforts we have made in the last 2 decades have been remarkably timid and poorly conceived.

III. Demand Subsidies

Unless one is careful and makes use of some simple economic theory, it is easy to fall into an all-too-common trap in discussions about innovation policy. The key point was signaled above in the switch from a discussion of spending on R&D to a discussion of the number of workers

engaged in this activity. To speed up growth, it is not enough to increase *spending* on research and development. Instead, an economy must increase the total *quantity of inputs* that go into the process of research and development. Spending is the product of a quantity and a price. To simplify the discussion, assume for now that people—scientists and engineers—are the key inputs in research and development. Formally, let E stand for spending on research and development and let N represent the number of scientists and engineers working in this area. Let w represent the average wage for these workers. Then trivially, $E = N \times w$. An increase in expenditure E will not necessarily translate into a corresponding increase in N, the number of scientists and engineers engaged in R&D. In principle, it is entirely possible for the entire increase in E to pass through as increases in w.

Continuing with the simplifying assumption that scientists and engineers are the only inputs in the production of research and development, we can illustrate how w is determined using the simple supply and demand framework presented in figure 7.3. The horizontal axis measures the number of scientists and engineers working in the private sector on R&D. The vertical axis measures their wage. The downward-sloping demand curve indicated by the solid line represents the


Figure 7.3
Wage w versus number of scientists and engineers working in private R&D

private return captured by a firm that hires some additional scientific workers and undertakes more research and development.

In the figure, the supply of scientists and engineers is represented by a vertical supply curve. The vertical slope of the curve reflects an assumption that the number of young people who become scientists and engineers and go to work in the private sector does not adjust in response to an increase in the wage that they receive for employment in the private sector. Section IV below outlines a much more complicated picture of the supply response of our educational system, but it is useful to start with the simple case of zero supply elasticity. To motivate this assumption, it is enough to keep the story from the preface in mind. The lack of information that is available to students who are making decisions about careers in science and technology suggests that our existing educational institutions may not lead to the kind of equilibration that we take for granted in many other contexts. If students do not have information about what wages will be, it will be much harder for them to adjust their career decisions in response to wage changes.

The downward-sloping dashed line in the figure represents the private demand for research workers when the government provides a subsidy for R&D. This subsidy could take the form of special tax advantages such as those afforded by the research and experimentation tax credit offered in the United States. Alternatively, the subsidy could take the form of cash payments to some firms as part of a cost-sharing agreement in which the government pays part of the cost of a research and development program. This is the kind of subsidy offered by partnership programs such as the Small Business Innovative Research (SBIR) grant program administered by the Small Business Administration or the Advanced Technology Program (ATP) administered by the Department of Commerce. Whether it comes in the form of tax credits or research grants, the effect of government spending is to shift up and to the right the demand for scientists and engineers who can perform the R&D.

From the perspective of a single firm, it seems obvious that a special tax incentive or a research grant will encourage the firm to hire more scientists and engineers and thereby to cause more inputs to be devoted to R&D. Yet one of the most basic insights in economics is that for the economy as a whole, things have to add up. If the total number of scientists and engineers is fixed, it is arithmetically impossible for employment of scientists and engineers to increase at all firms. As illustrated in the figure, if the supply curve of scientists and engineers is fixed, then the increase in demand induced by the subsidy will trans-

late into a proportional increase in wages for scientists and engineers with no increase in the inputs that are devoted to R&D.

It is important to recognize that this argument is separate from the usual concerns about "additionality" that have been raised with respect to R&D demand subsidy programs. People who focus on this problem are worried about how much the demand curve shifts. That is, they are worried that an additional dollar in subsidies does not translate into much additional private spending on R&D. This is a nontrivial issue. The evidence does seem to suggest that more generous tax treatment for R&D leads to higher *reported* levels of spending on R&D at firms. (See for example Bronwyn Hall and John van Reenen 1999.) An additional dollar in tax benefits seems to lead to about one additional dollar in reported R&D expenditure by firms. However, there is much less evidence about the extent to which this increase in reported R&D spending represents a true increase in spending relative to that which would have taken place in the absence of the credit. It is quite possible that some of this spending comes from relabeling of spending that would have taken place anyway. Deciding what qualifies for this credit is apparently a nontrivial problem for the tax authorities. Between 20% and 30% of claimed expenditures by firms are disallowed each year (National Science Board 1998:4-48).

For the SBIR program, Josh Lerner (1999) finds that firms that receive grants from the government experience more rapid sales and employment growth than a comparison group of firms selected to be similar to the recipient firms. This could be an indication that firms that receive grants do devote more inputs to R&D. But it could also reflect unobserved, intrinsic differences between the control group, which was constructed ex post by the researcher, and the recipient group, which was selected on the basis of a detailed application process that was designed to select particularly promising firms. In related work, Scott Wallsten (2000) finds that firms that receive a research grant from the government under the SBIR program seem to substitute these grant funds for other sources of funds, with little or no net increase in spending on R&D.

For both the tax credit and direct grant programs, we can identify a coefficient m which measures the true increase in private spending on R&D associated with each additional subsidy dollar from the government. In each case, there is some uncertainty and debate about how large this coefficient is. But for any positive value of m, the argument outlined above shows that the entire increase in spending may show up as higher wages for the existing stock of workers, with no increase

in the actual quantity of research and development that is performed. As a result, even a well-designed and carefully implemented subsidy could end up having no positive effect on the trend rate of growth for the nation as a whole.

Recent work by Austan Goolsbee (1998) suggests that, at least in the short run, the wage changes implied by a weak supply response are apparent in the data. He compares census data on wages for research workers with time series data that capture the variation in government spending on R&D. Direct government spending is well suited for this kind of analysis because it does not suffer from the concerns about additionality that are present for government subsidies for R&D. Surprisingly, using only these crude data, he finds strong effects on wages. For example, during the defense build-up between 1980 and 1984, federal spending on R&D increased, as a fraction of GDP, by 11%. His estimates suggest that this increased wages for physicists by 6.2% and aeronautical engineers by 5%.

In the face of this argument, defenders of demand-side R&D subsidies can respond in three ways. First, they can argue that people are not the only inputs used in R&D. If other inputs such as computers and specialized types of laboratory equipment are supplied elastically, then government subsidies for R&D could increase the utilization of these other inputs even if the number of scientists and engineers remains constant. If this were truly the intent of the various subsidy programs, it would be much more cost-effective for the government to provide the subsidies directly for these other inputs. Salaries account for the majority of total R&D spending. For example, in university-based research, annual research expenditures on equipment during the last decade have varied between 5% and 7% of total research expenditures (National Science Board 1998). If the goal of the subsidy program were to increase the equipment intensity of research and development and if the ratio of spending on equipment in the private sector is comparable to the figure for universities, a special tax subsidy for the purchase of equipment used in research would be substantially less costly than one that is based on total expenditures including salaries. Similarly, the government could achieve substantial savings, and still increase the use of equipment in R&D, if it restricted the grants provided by the SBIR and ATP programs so that these funds could be used only for additional purchases of equipment.

In the case of the targeted grant programs administered by the ATP or the SBIR, a defender could argue that even if the existing research subsidies do not increase employment of scientists and engineers in the

economy as a whole, they can increase employment at the recipient firms, at the cost of a reduction in employment at other firms. If government agencies were able to identify an allocation across firms and projects that is better than the one the market would implement, the targeted grant programs could still be socially valuable. But even the strongest supporters of the subsidy programs are hesitant to make this kind of claim about the superiority of government allocation processes. Note also that because the research and experimentation tax credit is available to all firms, it cannot be justified on this kind of basis of any hypothesized ability of the government to improve the allocation of research inputs between firms and projects.

If the goal is not to encourage equipment investment in the R&D sector or to give the government a bigger role in deciding how to allocate scarce R&D personnel, some other motivation must lie behind these spending programs. The final response could be for a defender of these programs to dispute the basic assumption behind the supply-and-demand model outlined here and argue that, at least in the long run, the supply of scientists and engineers working in R&D in the private sector does respond to demand-induced changes in wage. But to make this case, one must confront some of the peculiar features of the educational system that actually produces these highly skilled workers and ask if there are more cost-effective ways to increase the supply of these types of workers.

IV.   Overview of the Supply of Scientists and Engineers

Figures 7.4 and 7.5 give a broad overview of trends in the supply of scientists and engineers in the United States. Figure 7.4 updates data presented by Chad Jones (1995) on the number of scientists and engineers in the United States who are employed in research and development. These data are scaled by the size of the labor force. They show an increase in R&D employment as a fraction of the labor force from about 0.3% of the labor force in 1950 up to about 0.8% in the late 1960s, with no strong trend thereafter. The underlying data for this figure are collected by the NSF. (Data since 1988 are taken from Table 3-15 from National Science Board 1998.)

Official statistics on formal research and development capture only part of the private sector effort directed at innovation. Also, no consistent data series on employment in R&D is available in years prior to 1950. To give a more comprehensive overview of the proportions of skilled workers in the labor force over a longer time horizon, figure 7.5

## Scientists and Engineers in R&D as a Fraction of the Labor Force

Figure 7.4
Scientists and engineers in R&D as a fraction of the labor force
Data from Chad Jones 1995.

Figure 7.5
Engineers as a fraction of the labor force
Data are from Historical Statistics of the United States, 1975, and Statistical Abstract of
the United States, various years.

presents data on the total number of engineers as a fraction of the labor force, using occupational data collected by the Bureau of the Census. This series shows a similar pattern. Engineers increase, as a fraction of all workers, from the turn of the century up until 1970, and remain roughly constant thereafter.

Taken together, these figures offer little reassurance that the aggregate supply of scientists and engineers responds efficiently to market demand. Of course, it is logically possible that the growth in the demand for scientists and engineers experienced a sharp fall starting in the late 1960s. However, other labor market evidence based on relative wages such as that presented by Katz and Murphy (1992) suggests that a process of skill-biased technological change that raised wages for skilled relative to unskilled workers continued at about the same pace in the 1970s and 1980s as in the 1960s. Other work (see for example, Autor, Katz, and Krueger 1998) suggests that, if anything, the rate of skill-biased technological change actually increased in the period from 1970 to 1995 relative to the period from 1940 to 1970. Taken together, these data on quantities plus the independent evidence on the demand for skill suggest that one look more carefully at other possible factors that could influence the supply of scientists and engineers.

Figure 7.6 gives a schematic outline of the process that actually determines the supply of scientists and engineers. The two key stages in the production process are undergraduate education and graduate education. (For simplicity, graduate programs that lead to a terminal master's degree are grouped in this figure with those that provide Ph.D. level training.) The first major branch in the process distinguishes undergraduates who receive degrees in the natural sciences or engineering (NSE degrees) from those who receive all other types of degrees. Section V below looks at the possible nonmarket forces that could constrain this decision. After a student receives an undergraduate NSE degree, she can either go to work in the private sector or continue on to receive graduate training. Section VI looks at recent developments in the market for people with an advanced degree in the natural sciences or engineering.

V.   The Supply of Undergraduate Degrees in Science and Engineering

The market for education suffers from pervasive problems of incomplete information. Students contemplating a choice between different institutions typically have very little information about the
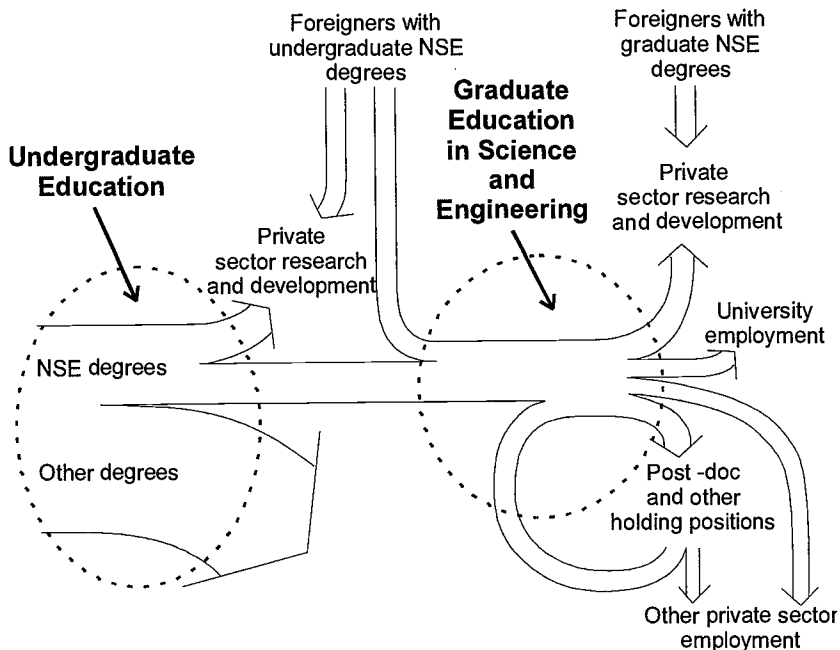
Figure 7.6
Schematic outline of the process determining the supply of scientists and engineers.

value-added they can expect to receive from one institution versus another. Employers selecting among graduates from different institutions also have very little objective basis for judging the absolute achievement levels of students from different schools, or even about students from the same school who have followed different courses of study. The competitive strategy that seems to have emerged in this market is one where undergraduate institutions have developed extensive systems for screening students by ability level. They enroll the most able students that they can attract. The schools compete for these students in large part by publicizing the degree of selectivity. Students, in turn, compete for admission to the most selective institutions because a degree from a more selective institution offers a stronger signal about the student's ability. Using data from the different campuses of the University of California, Robert Frank and Philip Cook (1995) suggest that competition along these lines has been getting more intense. For example, over time, SAT scores for students attending Berkeley, the UC campus that is perceived to be the most selective, have been increasing relative to SAT scores at other campuses.

In this kind of competitive environment, the traditional liberal arts university may face little pressure to respond to changing market demands for different types of skills. For example, imagine that government subsidies increased the market wage for scientists with several years of training beyond the undergraduate degree. Imagine that students are somehow informed about this change in the wage and respond by increasingly enrolling in undergraduate science courses that will prepare them for further study in engineering and science. A liberal arts university that has a fixed investment in faculty who teach in areas outside of the sciences and that faces internal political pressures to maintain the relative sizes of different departments may respond to this pressure by making it more difficult for students to complete a degree in science. Faculty in the departments that teach the basic science courses will be happy to "keep professional standards high" and thereby keep teaching loads down. Faculty in other departments will be happy to make study in their departments more attractive, for example by inflating the average grade given in their courses.

There is clear evidence that this kind of response currently operates on campuses in the United States. First, the number of students who begin their undergraduate careers with the intent of receiving a degree in science and engineering is substantially higher than the number who actually receive such a degree. For example, for white students, 12% of entering students intend to major in natural sciences and 9% plan to major in engineering. Only 8% of graduating students actually receive a degree in natural sciences and only 5% receive a degree in engineering (National Science Board 1998). For minority students, the attrition rate is even higher.

One additional indication of the pressure to shift students out of science and engineering degrees comes from the difference in the distributions of grades offered in courses required for degrees in these areas as opposed to grades in other courses of study. Measuring this difference is not straightforward because even within a department such as mathematics, and even within a specific subject area such as linear algebra, there are courses with easier grade distributions that are intended for people who will not continue toward a degree in science, and courses with a lower distribution of grades for people who will.

For example, students who place out of the basic calculus course on the basis of an advanced placement exam are more likely to take more difficult math courses than students who do not. This tends to lower the average grade they receive in the second-level math courses that they take. If one does not correct for this fact, one finds that math

grades for the students who place out of calculus are not, on average, any higher than math grades for students who do not place out of calculus. However, if one holds constant the specific second-level math courses that students take and compares grades for students with different backgrounds who take the same course, it is clear that students who have placed out of calculus do receive higher math grades than other students taking the same class (Rick Morgan and Len Ramist 1998).

To do this kind of analysis, the College Board, which administers the advanced placement exam, collected data from a representative sample of 21 selective universities. Using these data, one can do a direct comparison of grade distributions across different fields of study. Take, for example, the sample of second-level math courses that students who place out of calculus attend. These tend to be biased toward the classes that students majoring in mathematics or the natural sciences will take. One can then compare the distribution of grades in these courses with the distribution of grades in second-level English courses taken by students who receive advanced placement credit in English composition; or with the distribution of grades in second-level history courses taken by students who receive advanced placement credit in American history. As table 7.1 shows, in the selected math courses, 54% of all students received a grade of A or B. For the English courses, the fraction with an A or a B is 85%. For the history courses, the fraction is 80%. For social science courses such as political science or economics, the fraction of students who receive a grade of either A or B is about 75%.

As figure 7.6 shows, immigration is an alternative source of supply for the labor market in the United States. If the domestic supply of scientists and engineers is constrained to a significant extent by our existing system of undergraduate education, one should see evidence that the response in terms of undergraduate NSE degrees differs from that of immigrants. Recently, much of the discussion of migration has focused on political pressure from technology-intensive firms for increases in the number of H1B visas that permit private firms to hire skilled workers from abroad to fill entry-level jobs in areas such as computer programming. This debate has obscured the extremely important role that immigration has long played in supplying scientists and engineers with the highest levels of skill. Moreover, immigration is clearly responsive to demand conditions. Fields such as computer science and engineering, where indicators suggest that market demand is high relative to the available supply, are the ones that have experienced

Table 7.1
Fractions of students receiving an A or a B in different subjects

Subject Area	Fraction of Students Receiving a Grade of A or B
English	85
History	80
Economics and Political Science	75
Mathematics	54

the largest inflows from abroad. For example, in 1993, 40% of the people in the United States who had a Ph.D. degree in engineering were foreign-born. In computer science, 39% of the Ph.D. holders were foreign-born. In the social sciences, where demand for new Ph.D. recipients is generally much lower (economics being a notable exception), only 13% of the Ph.D. holders were foreign-born (National Science Board 1998). These immigration flows stand in sharp contrast to the trends in undergraduate education. From the mid-1980s until 1995, the number of undergraduate degrees in engineering and in mathematics and computer sciences fell substantially. For example, in the 1980s and 1990s, as the personal computer and internet revolutions were unfolding, the number of undergraduate degrees in computer science showed no strong trend, increasing at first in the mid 1980s, then falling in the 1990s and ending at about the level at which it started in the early 1980s.

Engineering degrees follow a very similar pattern (National Science Board 1998: Table 2-20). Between 1981 and 1995 there is no change in the number of undergraduates who receive degrees in engineering. The number does increase in the late 1980s but then returns to the previous level. For future reference, note that the number of master's degrees in engineering behaves quite differently. From 1981 to 1995 the total number of master's degrees in engineering increased steadily so that the number in 1995 was about 1.7 times the number in 1981 (National Science Board 1998: Table 2-27).

Another sign that the domestic enrollment of students who are able to continue in science and engineering is a critical bottleneck comes from an examination of downstream developments in Ph.D. education. From the mid-1970s to the mid-1980s, the number of Ph.D. degrees awarded in the United States each year in natural sciences and engineering remained roughly constant at about 12,500. (Here, as elsewhere, natural sciences and engineering exclude behavioral and social

sciences.) Then, starting in 1986, this number began a steady increase up to 19,000 per year in 1995.

We can use this expansion in the size of Ph.D. programs to gauge the elasticity of the foreign supply response and compare it to the domestic supply response. In 1986, U.S. citizens accounted for about 8,000 of Ph.D. degree recipients, and noncitizens accounted for the other 4,500. In 1995, the number of degrees for U.S. citizens had increased by about 20% to around 10,000 and the number of degrees awarded to noncitizens had more than doubled to 9,000 (National Science Board 1998: Table 2-35).

A similar, though less extreme picture emerges from an examination of master's degrees, particularly in the high-demand areas of computer science and engineering. As market opportunities for holders of the master's degree increased and universities added to the number of slots that they made available in master's degree programs, foreign students responded more strongly than U.S. citizens, just as they did when new positions in Ph.D. programs opened up. In 1975, foreign students received 22% of the master's degrees in engineering and 11% of the master's degrees in math and computer science. By 1995, foreign students accounted for 39% of the master's degrees in engineering and 35% of the master's degrees in math and computer science (National Science Board 1998). In both instances, increased downstream demand for undergraduates with NSE degrees does not seem to have induced a sufficient supply response. The system equilibrated by importing more foreigners.

VI.   The Supply of Ph.D. Degrees in Science and Engineering

The sharp increase in the 1990s in the number of Ph.D.s granted has been accompanied by generally declining job prospects for degree recipients. In the most recent period, it is possible that part of the reason why undergraduate students did not pursue degrees in the natural sciences is that they were vaguely aware of the worsening job prospects that Ph.D. recipients faced. Note, however, that developments in the academic market for Ph.D.s cannot explain the absence of an increase in undergraduate degrees in engineering or in specialized areas such as computer science where job prospects for Ph.D. recipients have remained strong. Also, the weak market for new Ph.D.s would only have been a factor fairly recently, primarily since 1990 when the increased

supply of Ph.D. recipients began to show up on the market. Nevertheless, going forward, the weak academic market for some types of Ph.D.s will certainly be a complicating factor in any attempt at increasing the number of undergraduate degrees that are awarded in natural science and engineering. To increase the number of undergraduates who receive an undergraduate degree in the natural sciences and engineering, they must be convinced that this kind of degree can lead to better career outcomes than the dead-end postdoctoral positions that have become increasingly common in some fields.

Independent of its role in influencing undergraduate degrees in the United States, understanding the behavior of the market for Ph.D.s is critical to the formulation of policy concerning the supply of scientists and engineers. The thrust of the possible programs outlined below is to substantially increase this supply. Yet many people in the academic community are convinced that the most pressing science policy issue in the United States is the Ph.D. glut. They have advocated measures that would reduce the supply of Ph.D.-level scientists and engineers. A careful look at the market for Ph.D.s is necessary to explain why increases in the supply of scientists and engineers with several years of graduate training are still called for even in the face of difficulties in the labor market for Ph.D.s. The key point here is to distinguish between people who are trained exclusively for employment in research universities and people who can work in research and development in the private sector.

Look again at figure 7.6. Events in the Ph.D. market can be summarized in terms of this figure. As noted above, the total flow of students through NSE Ph.D. programs increased starting in the late 1980s and continuing through the 1990s. Much of this flow has been directed at two of the alternatives upon leaving graduate school—university employment and postdoc and other holding positions. The challenge in this area is not to increase the total numbers of Ph.D. degree recipients, but to increase the fraction of them that can put their skills to work in private sector research and development.

This pattern of outcomes—increased numbers of Ph.D. recipients and steadily worsening academic job prospects—can be explained by increased subsidies for Ph.D. training. These subsidies derived from increased support for university-based research, which is complementary to Ph.D. training. As a result, the nature of the support for graduate students changed along with the level. Consider the sample

of students who received their primary means of support for their Ph.D. education from the federal government. Between 1980 and 1995, the fraction whose primary mechanism of support was a traineeship fell from 25% to 15% and the fraction whose primary mechanism was a research assistantship increased from 55% to 63%. The fraction receiving their primary support from fellowships stayed roughly constant at about 10%. Among students whose primary support was from sources other than the federal government (primarily state governments), research assistantships also increased by about 10 percentage points (National Science Board 1998: Chapter 5).

Because this increase in supply consisted of people who planned to pursue academic research appointments, the increased supply of Ph.D. recipients was accompanied by generally worsening job prospects for Ph.D. recipients in the academic market. For example, consider in any year the sample of people with degrees in the natural sciences and engineering who were working in academic institutions and who had received their Ph.D. degree within the previous 3 years. In the early 1980s, there were about 17,000 of these recent degree recipients working in academic institutions. About half of them had faculty jobs. The rest held postdoctoral positions or some other form of appointment. By 1995, this same measure of recent Ph.D.s in academic institutions had increased to 23,000, but the number holding faculty positions remained roughly constant, at about 8,500. The entire increase of 6,000 recent degree recipients is accounted for by increases in nonfaculty appointments (National Science Board 1998: Table 5-29).

The problems in the academic market in the life sciences were documented in a report from the National Research Council (1998). In the last decade, this is the area that has benefited from the most rapid rate of growth of federal research support. Between 1970 and 1997, the median time to receipt of a degree increased by 2 years to a total of 8 years. The number of people who hold a postdoctoral appointment 3 or 4 years after receipt of the Ph.D. increased from 6% to 29% between 1973 and 1995. The fraction of Ph.D. recipients who do not hold a permanent full-time job in science and engineering 5 or 6 years after they have received their degree increased from 11% in 1973 to 39% in 1995. The 1995 data, which were the most recent available at the time that the National Research Council wrote its report, reflect long-term outcomes for the 1989–90 cohort of Ph.D. recipients. Because of the steady increase in the number of degree recipients throughout the 1990s, the competitive pressures in this field have probably worsened still further.

VII.   An Interpretation of the Evidence Concerning Higher Education

The picture that emerges from this evidence is one dominated by undergraduate institutions that are a critical bottleneck in the training of scientists and engineers, and by graduate schools that produce people trained only for employment in academic institutions as a side effect of the production of basic research results. This description of the system as a whole hides a heterogeneous mix of different types of institutions. Not all of them will behave according to the description given above.

For example, the pressure to keep enrollments down in the natural sciences and engineering will not be present at institutions that specialize in this kind of training. They may therefore face different kinds of incentives and behave differently in the competition for students. The institution that my son attends, Harvey Mudd College, is one of these specialized institutions, and this may explain why it features information about the market outcomes for its graduates more prominently than traditional liberal arts universities. A quick check of data from other schools is consistent with this observation. MIT and Caltech, two selective schools that also concentrate in science and engineering, present information about median salaries for their undergraduates on the web pages that provide information for potential applicants. Harvard and Stanford, two comparably selective institutions that cover the whole range of academic disciplines, apparently offer no information on their web pages about salaries or enrollments in graduate school.

One natural question that the model outlined here does not address is why competition by entry of more schools like Harvey Mudd, MIT, and Caltech has not partially solved the bottleneck problem described here. Mudd, which is about 50 years old, is a relatively recent entrant in this market, but in general, entry seems to be a relatively small factor in the competition between undergraduate institutions. Presumably the incomplete information available to students and employers about the quality of the education actually provided at any institution is a big factor limiting the entry process, but the nature of competition between schools deserves more careful consideration.

There are also different types of institutions that provide graduate education. The description offered here focuses primarily on graduate education in the sciences, which takes place almost exclusively within institutions where the revenue and prestige associated with research are more important motivating forces than tuition revenue. Training in

these departments differs sharply from the kind of training offered by professional schools where income from tuition is a much more important determinant of institutional incentives. It should not be surprising that, as my research assistant discovered, business schools and law schools follow very different strategies from the ones used by departments of science when they compete for students. In many ways, master's level training in engineering is like these professional schools. Much of the income associated with these programs comes from tuition. Departments that get to keep a portion of this master's level, but not of undergraduate tuition revenue, should therefore be willing to expand the size of their master's programs at the same time that they put limits on the size of their undergraduate programs. These kinds of incentive effects may help explain why master's degree programs in engineering have shown steady growth while undergraduate engineering degrees have not.

In its report on career prospects in the life sciences, the National Academy Board on Biology concluded that policymakers should restrain the rate of growth of graduate students in the life sciences. In my language (not theirs) they also recommended that graduate education in the life sciences be reshaped along lines that are closer to those followed by professional schools. They recommended that students be given more information about career prospects, that they be given training that prepares them for employment in jobs outside of university-based research, and that funds that support the training of graduate students be shifted away from research assistantships and toward training grants or other forms of support that give more control over a student's education to the student.

This last and most controversial recommendation is the one that has the greatest potential to shift the traditional science-based model of graduate education closer to the model that we see in master's level professional schools of tuition-paying customers who collectively can exert a significant degree of control over what happens during the process of education. Similar proposals for modifying Ph.D. training have been made by a variety of study panels. All have received mixed support at best from the scientific community as a whole. (See the discussion of this point in National Science Board 1998: 5-33.)

Opposition to any change in the form of support for graduate education is usually justified in public on the grounds that there is insufficient evidence about what the effects might be for any change in the system of funding for graduate students. A more fundamental

problem—one that goes largely unreported in print but that prominent scientists are willing to justify in private—is that the current funding and training system, one that puts graduate students in the position of apprentices to established scientists and that does not prepare students for careers outside of science, is crucial to the maintenance of the institutions of academic science. Recent work by Scott Stern (1999) offers convincing evidence that recipients of Ph.D. degrees exhibit a strong preference for engaging in the activities in science and are willing to accept substantial wage reductions if doing so will allow them to continue to pursue these activities. This preference could be the result of a selection process that attracts people with this taste into Ph.D. training in science, a training process that cultivates this taste, or a combination of the two. Regardless of the mechanism, any attempt to make the training of Ph.D. students resemble more closely the training of students in business schools could have the effect of significantly undermining the commitment to the ideas and process of science that Stern is able to document. This commitment, which may be psychologically and functionally similar to the commitment induced by training for membership in a religious order or a military unit, may be critical to the preservation of the institutions of science. Unfortunately, it may also help explain why the existing system of graduate education seems so poorly suited to training people for employment outside of academic science. For this combination of reasons, the task of modifying the educational system that trains scientists and engineers may be both very important and very delicate.

VIII.   Goals and Programs

To formulate growth policy, policymakers may want to start by distinguishing goals from programs. Goals should be conservative. They should represent objectives that are neither risky nor radical and for which there is a broad base of intellectual and political support. Goals should remain relatively constant over time. They should also imply metrics for measuring success. By these criteria, increasing the long-run trend rate of growth is not specific enough to be a goal. It is appropriately conservative and should be the subject of a broad consensus, but because it is so difficult to measure the trend rate of growth, it does not imply any workable metrics that we can use to measure progress toward the goal. In contrast, increasing the fraction of young people in the United States who receive undergraduate degrees in

science and engineering could qualify as a goal. So could increasing the total quantity of resources that are devoted to research and development.

In contrast to a goal, a program is a specific policy proposal that seeks to move the economy toward a specific goal. For example, the Research and Experimentation tax credit is a specific program that is designed to achieve the goal of increasing the resources used in research and development. It should be possible to judge the success of a program against the metric implied by the goal that it serves. All programs should be designed so that they can be evaluated on a policy-relevant time horizon. If they are, they can also be less conservative and more experimental than the underlying goals. A variety of programs could be tried, including ones where there is some uncertainty about whether they will succeed. If the evidence shows that they do not work, they can be modified or stopped.

To illustrate how this framework could facilitate better analysis of the growth process, it helps to focus on a specific set of hypothetical goals. Imagine that policymakers and the public at large accepted the following goals because they want to increase the long-run rate of growth in the United States. (1) Increase the fraction of 24-year-old citizens of the United States who receive an undergraduate degree in the natural sciences and engineering from the current level of 5.4% up to 8% by the year 2010 and to 10% by 2020. (2) Encourage innovation in the graduate training programs in natural science and engineering. (3) Preserve the strengths of the existing institutions of science. (4) Redress the imbalance between federal government subsidies for the demand and supply of scientists and engineers available to work in the private sector.

Each of these goals suggests natural metrics for measuring progress. The NSF currently measures the fraction of 24-year-olds who receive undergraduate degrees in the natural sciences and engineering (NSE). These data are also available for other countries. Although the United States provides undergraduate degrees to a larger fraction of its young people than almost all other developed nations, many fewer undergraduates in the U.S. receive degrees in natural science and engineering. As a result, the fraction of all 24-year-olds with undergraduate NSE degrees is now higher in several nations than it is in the U.S. The United Kingdom (8.5%), South Korea (7.6%), Japan (6.4%), Taiwan (6.4%), and Germany (5.8%) all achieve levels higher than the 5.4%

level attained in the United States (National Science Board 1998: Chapter 2). The experience in the United Kingdom also shows that it is possible to expand this fraction relatively rapidly over time. In 1975, the figure there stood at only 2.9%.

The indicators for the next two goals will have to be more eclectic. Possible indicators of innovation in graduate education could include the creation of graduate training programs in new areas (bioinformatics, for example) where the private sector demand for graduates is high; or programs that involve new types of training (internships in private firms, perhaps); or programs that offer different types of degrees from the traditional master's or Ph.D. One would also like to see continued strength in the Ph.D. programs that form the core of our system of basic scientific research, measured perhaps by the quality of students that they attract both domestically and from abroad. The second and third goals explicitly allow for the possibility that developments in these two areas need not be closely linked. Universities might introduce new programs in an area such as bioinformatics that train people primarily for work in the private sector without affecting existing programs in biology. The new programs could have the same independence from Ph.D. training in biology that programs of chemical engineering have from Ph.D. training in chemistry. As a result, innovation in the sense of new programs need not imply any changes in the existing Ph.D. training programs and need not take any funding from those programs. If the country makes progress toward the first goal, and the number of U.S. citizens who pursue undergraduate studies in science increases, this could improve the quality of the domestic applicant pool for the traditional Ph.D. programs at the same time that it supplies people to the new alternative forms of graduate education.

It will take new funding from the federal government to encourage the introduction of new training programs and still preserve the strength of existing graduate programs. The last goal sets a rough benchmark that policymakers might use to set expectations for how much funding might be allocated on a permanent basis toward these goals. In the last 2 decades, the primary programs that have subsidized the private sector demand for R&D have been the research and experimentation tax credit, the SBIR program, and the ATP program. Rough estimates of the costs for these programs are $1 billion each per year for the tax credit and the SBIR program and between $300 and $400 million

per year for the ATP program. The fourth goal suggests a starting target of around $2–2.5 billion per year in subsidies for the supply of scientists and engineers.

If policymakers adopted these kinds of goals, then it would be a straightforward process to design programs that might help achieve them and to evaluate these programs after they are implemented. The following list of programs only begins to suggest the range of possibilities that could be considered. (1) Provide training grants to undergraduate institutions that are designed to increase the fraction of students receiving NSE degrees. (2) Finance the creation of a system of objective, achievement-based (rather than normed) tests that measure undergraduate level mastery of various areas of natural science and engineering. (3) Create and fund a new class of portable fellowships, offered to promising young students, that pay $20,000 per year for 3 years of graduate training in natural science and engineering.

The details for all of these programs would have to be adjusted based on more detailed prior analysis and as experience with any of them is acquired in practice. Many alternative programs could also be proposed. These three are offered here primarily to indicate the wide range of possibilities and to move the debate about government programs out of the rut in which it has been stuck for some 20 years.

Training grants could be very flexible. They could follow the pattern that has already been established for training grants at the graduate level. Formally, grants could still be given to a lead principal investigator, but in effect, they would offer financial support to a department at a university or college. The details of the proposed training program would be left open to the applicants. Like all grants, they would be peer reviewed, with fixed terms but renewable. One of the central criteria in evaluating any proposed grant would be some estimates of its cost-effectiveness as measured by the expenditure per additional undergraduate NSE degree granted. At this point, undergraduate institutions in the United States award about 200,000 NSE degrees each year. The vast majority (roughly 95%) of these degrees are awarded to U.S. citizens. It will take an increase of about 100,000 NSE degrees to U.S. citizens per year to meet the goal of having 8% of 24-year-olds receive an NSE degree. If the federal government devoted $1 billion per year, or about $10,000 per additional degree recipient as a reward to schools that could increase the numbers of NSE degrees that they award, universities would surely find it in their interest to reverse the existing pattern of discouraging students from pursuing NSE degrees. Existing

liberal arts universities could reallocate resources internally. Specialized science and engineering schools could use these funds to expand. New institutions could enter the educational marketplace.

One of the obvious risks associated with a goal of increasing the number of NSE degrees is the risk that universities would simply relabel existing degrees as NSE degrees or would substantially reduce the content of the NSE degrees that they award. One additional criterion for evaluating training grants would be the presence of metrics that verify whether the quality of the degree from the recipient institution is being eroded. But eventually, it would be more efficient to have objective, national measures of student mastery of science rather than the kind of implicit, idiosyncratic, institution-specific assurances of the quality that universities now provide. The model for this system of measures would be the advanced placement tests offered to high school students by the College Board. This organization has shown that it is possible to construct reliable tests with the property that when teachers teach to the test, the students actually learn the material that they should learn. Just as the AP system is guided by high school and college educators, one would expect that any such system for measuring undergraduate achievement in science would be guided primarily by the professors who teach science at the undergraduate and graduate level. Presumably, scores on these kinds of achievement-based exams would not replace other indicators like course grades, letters of evaluation, and general measures of intellectual ability such as are provided by the existing graduate record exams. Nevertheless, they would provide a new and useful piece of information about performance by individual students, by different educational institutions, and by the nation's educational system as a whole. Given the pervasive problems of incomplete information in higher education, it would surely be of value to students, employers, and faculty members to have access to objective measures of what students actually learn.

The new fellowship program is intended primarily to encourage the process of innovation in graduate education by providing a ready pool of funds that could be spent on any attractive new programs that are created. It would also create additional incentives for students to pursue undergraduate NSE degrees. Possible details for such a program could be as follows: The government could select a sample of graduating high school students who show promise in science, say the more than 100,000 high school students per year who pass the advanced

placement exam in calculus. It could offer to a randomly selected treatment subgroup a fellowship that will pay $20,000 per year for 3 years of *graduate* education in natural science or engineering if the student receives an undergraduate NSE degree. (There would be little reason to pay them a subsidy for undergraduate education. Virtually all of these students already go on to get an undergraduate degree.) Granting the award before they begin their undergraduate study would allow them to take the science courses that prepare them for graduate study. Because the treatment group would be randomly selected, it will be easy to verify whether these grants increase the likelihood that a student receives an undergraduate NSE degree. One could also look among the students who continue their studies in graduate school and see whether the recipients of the portable fellowships select career paths that differ from the students who are supported under the existing RA and TA positions. To the extent that fairness is a concern, one could give some other award to the students in the control group, a new personal computer perhaps.

These fellowships would be portable both in the sense that they could be used to pay for training in any field of natural science and engineering and in the sense that they could be used at any institution that the student selects. Some of the students who receive these fellowships would no doubt pursue a traditional course of Ph.D. study, but some may be willing to experiment with other kinds of degrees. Because these funds would represent new funds, not subtractions from the funds that are already used to support graduate students, and because they would only cover 3 years of training, they should not pose much risk to the traditional training system in basic science.

If the government paid for a total of 50,000 of these fellowships each year, or about 16,700 for each annual cohort of students, this would represent an annual expenditure of about $1 billion. (To pay for 16,700 new fellowships each year, the government would presumably have to offer many more because the take-up rate would be less than 100%.) It is possible that the availability of these funds would not lead to the introduction of new courses of study that cater to the recipients. If this were the outcome, the fellowships would be judged a failure and would presumably be discontinued. But, a priori, it seems quite likely that a flow of funds of this magnitude would induce at least some innovative response from our educational system. It should not take many years of observation to verify whether this conjecture is correct.

IX.   Conclusions

The analysis here is driven by two basic observations. The first is that better growth policy could have implications for the quality of life in all dimensions that are so large that they are hard to comprehend. The second is that in the last several decades, the efforts that our nation has undertaken to encourage faster growth have been timid and poorly conceived.

We owe it to our children and their children to address questions about growth policy the way we would approach a major threat to public health. We must use the best available evidence and careful logical analysis to frame new initiatives. We must then be willing to run experiments and to see what actually works and what does not.

References

Autor, David, Lawrence F. Katz, and Alan B. Krueger. 1998, November. "Computing Inequality: Have Computers Changed the Labor Market?" *Quarterly Journal of Economics* 113:1169–213.

Frank, Robert, and Phillip Cook. 1995. *The Winner-Take-All Society*. New York: Free Press.

Goolsbee, Austan. 1998, May. "Does Government R&D Policy Mainly Benefit Scientists and Engineers?" *American Economic Review* 88:298–302.

Griliches, Zvi. 1992. "The Search for R&D Spillovers." *Scandinavian Journal of Economics* 94:347–520.

Hall, Bronwyn, and John van Reenen. 1999. "How Effective Are Fiscal Incentives for R&D? A Review of the Evidence." Working Paper no. W7098, National Bureau of Economic Research, Cambridge, MA.

Historical Statistics of the United States: Colonial Times to 1970, Part 1, U.S. Bureau of the Census, Washington, D.C., 1975.

Jones, Charles I. 1995, May. "Time Series Tests of Endogenous Growth Models." *Quarterly Journal of Economics* 110:495–525.

Jones, Charles I., and John C. Williams. 1998, November. "Measuring the Social Return to R&D." *Quarterly Journal of Economics* 113:1119–35.

Katz, Lawrence F., and Kevin M. Murphy. 1992, February. "Changes in Relative Wages, 1963–1987: Supply and Demand Factors." *Quarterly Journal of Economics* 107:35–78.

Langlois, Richard N. and Mowery, David C. 1996, "The Federal Government Role in the Development of the U.S. Software Industry." In David C. Mowery, ed., *The International Computer Software Industry: A Comparative Study of Industry Evolution and Structure.* New York: Oxford University Press.

Lerner, Josh. 1999, July. "The Government as Venture Capitalist: The Long-Run Impact of the SBIR Program." *Journal of Business* 72:285–318.

Maddison, Angus. 1995. *Monitoring the World Economy, 1820–1992.* Paris: Development Centre of the Organisation for Economic Cooperation and Development.

Morgan, Rick, and Len Ramist. 1998, February. "Advanced Placement Students in College: An Investigation of Course Grades at 21 Colleges." Report no. SR-98-13, Educational Testing Service, Princeton, NJ.

National Research Council. 1998. *Trends in the Early Careers of Life Scientists.* Washington, DC: National Academy Press.

National Science Board. 1998. *Science and Engineering Indicators—1998.* Arlington, VA: National Science Foundation.

Rosenberg, Nathan. 1969. *The American System of Manufactures: The Report of the Committee on the Machinery of the United States 1855, and the Special Reports of George Wallis and Joseph Whitworth 1854.* Edinburgh: Edinburgh University Press.

Statistical Abstract of the United States. U.S. Bureau of the Census, Washington, D.C., various years.

Stern, Scott. 1999, October. "Do Scientists Pay to Be Scientists?" Working Paper no. 7410, National Bureau of Economic Research, Cambridge, MA.

Wallsten, Scott. 2000, forthcoming, Spring. "The Effects of Government-Industry R&D Programs on Private R&D: The Case of the Small Business Innovation Research Program." *The Rand Journal of Economics.*