Technology and Growth: An Overview

Jeffrey C. Fuhrer and Jane Sneddon Little*

During the 1990s, the Federal Reserve has pursued its twin goals of price stability and steady employment growth with considerable success. But despite—or perhaps because of—this success, concerns about the pace of economic and productivity growth have attracted renewed attention. Many observers ruefully note that the average pace of GDP growth has remained below rates achieved in the 1960s and that a period of rapid investment in computers and other capital equipment has had disappointingly little impact on the productivity numbers. Others see faster growth as softening the impact of widening income inequality or the stagnant real wages earned by many citizens.

Most of the industrial world has experienced a similar decline in trend and productivity growth, an increase in income inequality, and even slower job creation than we have seen here in the United States. While some (particularly Asian) developing countries are rapidly joining the ranks of the industrialized, most remain mired in poverty. According to the World Bank's recent report on poverty, over 20 percent of the world's population lives on less than one dollar a day. This situation wastes human talent and contributes to political instability.

While raising trend growth rates would not directly address distributional issues, increasing growth rates by even a fraction of 1 percent would, with compounding, have profound implications. As Robert Lucas has pointed out, "the consequences for human welfare are simply staggering. Once one starts thinking about them, it is hard to think of anything else." Unfortunately, economists and policymakers do not

*Vice President and Economist, and Assistant Vice President and Economist, respectively, Federal Reserve Bank of Boston.

know how to engineer such an outcome. While the determinants of growth are widely agreed to be capital, labor, and a composite including managerial skills and organizational culture that Robert Solow abbreviated as "technology," the interrelationships among these variables are not clearly understood. In the developed economies, at least, recent large capital investments have shown surprisingly little positive impact on productivity or potential growth. Accordingly, attention has increasingly turned to the role of such intangibles as human capital, social organization, and technology.

Because these puzzles are so compelling, the last few years have seen a resurgence in research on the economics of growth. This groundswell reflects the availability of new data bases and an improved ability to model imperfectly competitive conditions. Primarily, however, this enthusiasm indicates that many members of the economics profession concur with *The Economist* (June 1, 1996) that "understanding growth is surely the most urgent task in economics." For these reasons, the Federal Reserve Bank of Boston devoted its fortieth economic conference, held in June 1996, to *Technology and Growth.* We hoped to explore what we know and clarify what we do not know about these issues.

A number of themes emerged from the discussions. For the most part these themes took the shape of questions repeated in various contexts. For example, one fundamental question asked throughout the meeting was just how important is technology—to growth, to productivity, to convergence? The answer, it was generally agreed, depends on one's definition of technology, with the majority favoring an inclusive approach. Most participants were sympathetic with the need to decompose technology into its constituent parts—innovation, development, and diffusion—and to include intangibles like organizational structure, management skills, and culture in the package labeled technology. Another theme that arose early on and reappeared throughout the conference was the unpredictable nature of technological change and the consequences of our uncertainty (or lack of imagination) concerning its ultimate path.

A third motif involved the role of innovation and the importance of knowledge-based spillovers within the growth process. While early work based on Robert Solow's model attributed most growth to exogenous technological change, more recent neoclassical research, exemplified by Dale Jorgenson's work, has greatly reduced technology's role by broadening our definition and improving our measures of capital. Indeed, Jorgenson concludes that human and physical capital accumulation, properly measured, explains almost all growth with little scope for innovation or knowledge-based spillovers.

But not everyone is fully persuaded that capital accumulation, however defined, can by itself account for the great bulk of welfare improvements experienced in recent decades. Noting a major inconsistency between the rate of convergence to steady state growth rates predicted by the neoclassical approach and the slower rate observed in fact, the new growth theorists give technological change, rather than capital, a bigger role in the growth process. They argue that technological change requires human effort and is, therefore, not exogenous, that the returns to R&D and other knowledge-based investments are not fully appropriable, and that spillovers from innovation have contributed importantly to growth. Naturally, thus, the new growth theorists stress the need to model the innovative process and the role played by these spillovers. While participants of both camps generally favored developing fully endogenous models, they disagreed about our current or potential ability to meet this challenge and, more basically, about its actual importance. In this regard, most, but not all, of the participants believe that spillovers are pervasive and significant.

A further theme was the need to be realistic in at least two areas. First, we need to acknowledge that potential growth may not return to its pace in the 1960s and that we may have to be satisfied with raising the level of output rather than the rate of growth. Economists also need to admit how little we understand about the growth process and how small are the likely consequences of the policy measures we advocate.

The conferees did agree on several points. Since the previous heyday of growth economics in the late 1950s, economists have greatly improved their ability to model the growth process by broadening their definitions and measures of physical and human capital. This development has reduced the role of exogenous technological change and narrowed the differences between the neoclassical and new growth theorists. Remaining areas for dispute and research include the need for modeling the various components of technology and the interactions between the determinants of growth and the growth process itself. Moreover, although research has not clearly demonstrated that the technology embodied in widely available capital equipment has much impact on productivity, participants generally concurred that technology defined to include management, social organization, and culture is likely to be important.

As for policy recommendations, conference participants largely agreed that the path of technical development and diffusion is highly unpredictable. Given this uncertainty and the gap between the social and private returns to R&D, most participants favored modest and balanced public support of basic research and other pro-competitive policies. They were less convinced about the benefits of the patent system.

On the macro side, participants universally endorsed the need to reduce fiscal deficits in order to promote saving and investment and the desirability of maintaining open trading systems in order to spur innovation. Several attendees advocated greater use of consumption-based tax systems. Many also saw an ongoing need for government investment in education and training, in limited amounts of R&D, and in improved statistical capabilities. Monetary policy's contribution was generally seen to be limited to maintaining price stability, but Bob Solow reminded us that balancing relatively tight fiscal policy with relatively accommodative monetary policy tends to favor growth. He also noted that below-potential growth discourages investment and innovation. Finally, if increasing productivity growth remains out of reach, some participants saw a need for more generous redistributive policies.

Keynote Address: The Networked Bank

In his keynote address, **Robert M. Howe** provided an intriguing view of how one industry—financial services—has responded to rapid technological change, and a vision of how that industry will be transformed with the introduction of technologies already in the development pipeline. Howe's vision is that of the networked economy: "the integration of people and institutions obtaining information, transacting business, and entertaining and educating themselves in a connected world, with electronic networks as the underlying backbone." In addition to detailing the modifications required of banks to survive in this networked environment, Howe shows where consumers fit into this system.

The networked bank has three components. The first component includes the access channels that link the consumer to the bank—ATMs, telephones, PCs, and bank tellers. Control over these channels rests in the hands of consumers and of third-party providers, such as on-line services. The second component is the "customer information and relationship management system," the bank's data base tracking customer activities to glean information about customer preferences. Howe suggests that effective use of this information—to tailor products to individual consumers or to determine the bank's most profitable market segments—will become the bank's "most valued asset and source for competitive advantage." The third component of the networked bank is the "core back-office system," which coordinates the operational systems, retail and commercial banking functions, and alliances with other service providers—for example, insurance firms or travel agents—that offer their services through the bank.

Howe forecasts the emergence and widespread distribution of a suite of new technologies that will support the networked bank. These include improved communications interfaces, such as speech and handwriting recognition; three-dimensional, high-resolution graphics; and touch screens. Network infrastructure will improve rapidly in speed and price, and user-screening and encryption will enhance security. In addition, the continued miniaturization of processor and storage technology will allow smart cards with PC capabilities for financial transactions, inventory control, or transmission of medical patient information. "Intelligent agent" software will respond to a consumer's complex queries; for example, "Go find me the lowest-priced Brand X automobile with the following features." Finally, networked banks will make greater use of new tools for data management, to analyze customers and transactions for targeted marketing campaigns.

These changes in the competitive environment pose new challenges to banks. Because the provision of a service will often involve a number of players, banks must establish "electronic value chains" that link the bank, the customer, perhaps a vendor, and a network infrastructure provider. Howe foresees a notable shift of power from banks to consumers and providers of access channels. With easy access to many options, a consumer may have little loyalty to a particular financial institution. A bank will need to differentiate its product from its easily accessible competitors, even when its product may appear only as a menu item on a screen. The bank's most valuable asset will shift from its branches, the current interface with its customers, to its customer data base and its expertise in extracting useful information from that data base.

How can the networked bank respond to these challenges? Howe proposes three possible strategies. The first, the "customer-centric" strategy, uses the bank's customer data base and data base analysis to serve each customer with unique, customized services. A second response is the "life-event" strategy: The bank becomes the provider of a cluster of services required by the consumer at key life events, such as buying a house or planning for retirement. A third option is the commodity strategy, in which the bank competes by providing standardized services through a wide range of access channels at the lowest cost.

Finally, Howe points out that these technological advances pose difficult questions for financial regulators. For example, does a global electronic financial system imply greater systemic risk to the payments system? How are standards of security and reliability established for new products? Will new clearinghouse organizations be required for new products? How are consumers to be protected if non-regulated industries can offer bank-like services? Who guards the consumer's right to privacy?

Technology in Growth Theory

Dale Jorgenson's paper traces the economics profession's understanding of technology and economic growth from the seminal works of Harrod (1939), Domar (1946), Solow (1956), and Kuznets (1971) to the more recent "endogenous growth models" of Grossman and Helpman (1994). In Jorgenson's view, the profession formed a rare and temporary consensus in the 1970s around the neoclassical growth model of Solow and the empirical work of Kuznets. Solow's simple theoretical framework, which decomposed contributions to output according to a constant-returns-to-scale production function with capital and labor as inputs, "provided conceptual clarity and sophistication." Kuznets' complementary work linking measures of capital and labor inputs to final output provided "persuasive empirical support" for the neoclassical growth model by documenting the correlation among inputs and outputs for the United States and 13 other developed countries over a long historical span. What stands out most for Jorgenson about these twin pillars of early growth theory, theoretical and empirical, is the lack of integration between them.

In early implementations of the Solow growth model, growth arose primarily as a result of increases in productivity. Because the reasons behind productivity increases were not understood, most economic growth was attributed to exogenous causes that largely reflected, as Abramovitz (1956) phrased it, a "measure of our ignorance." The contribution of investment in physical and human capital was assumed to be relatively minor.

Work by Jorgenson and others in the 1980s has attempted to diminish our ignorance by using carefully constructed measures of the inputs to production in an econometric model. The product of this research strategy is a model that fully characterizes the accumulation of human and physical capital and attributes almost all of economic growth to increases in the rate of capital accumulation, once properly measured. A truly satisfactory model of endogenous investment in new technology has eluded the profession thus far, however, in large part because of the difficulties inherent in measuring the output of the research and development sector (a problem first identified by Griliches in 1973).

Interest in growth theory waned in the 1970s, in the aftermath of the oil price shocks and a renewed attention to the determinants of business cycle fluctuations, but the debate over "convergence" in the 1980s and early 1990s revived interest, even as it challenged the validity of the Solow framework. Because the convergence debate focused on the long-run growth experience of nations, it brought to light a key question that had not previously been addressed: Could private investment, whose returns accrue only to the investor, account for the leaps and bounds in output that some countries have observed over centuries? Or do we need "spillovers" in "knowledge capital," which may result from individuals' investment but which benefit all, to explain growth over long spans of time?

Jorgenson describes the essence of the convergence debate as follows: If Solow's model is approximately correct, then over a long enough period of time, a country will converge to its "steady state" or long-run rate of per capita income growth, which is determined by its saving and population growth rates. The Solow model predicts that the rate of convergence to the steady state will depend upon the share of capital in GDP, the rate of population growth, the rate of productivity growth, and the rate of depreciation of capital equipment. Using plausible estimates of these determinants for many countries implies a rate of convergence of about 4 percent per year. While empirical studies have found evidence of convergence, the estimated rate of convergence—about 2 percent per year—is too slow to be consistent with the Solow model.

An influential paper by Paul Romer (1986) highlights the inconsistency between the simple Solow model and the evidence on rates of convergence. Romer deduced that, for the slow observed rates of convergence to be consistent with the Solow model, the share of national income devoted to capital accumulation must be about twice as large as normally assumed. The reasoning is as follows: The larger is the share of national income devoted to capital accumulation, the more investment is required to increase output; the more investment is required, the slower will be the convergence to the steady state for a given investment rate.

Because doubling the share of income going to investment is just a "crazy explanation" of the slow-convergence puzzle, Romer and others suggest what they consider to be more plausible alterations to the standard growth model, such as increasing returns to scale in the aggregate production function, and spillovers of the returns to private investment to the rest of the economy. In their view, only these alterations can reconcile the standard growth model with the convergence data.

Mankiw, David Romer, and Weil (1992) find, however, that Paul Romer's crazy explanation is unnecessary and the Solow model can be resurrected once one controls for differences in human capital across countries. Allowing for these differences again reconciles the basic Solow model with the share of capital in the value of output and with the slow rates of convergence observed over time across countries.

A recent paper by Islam (1995) extends this work, allowing for different levels of productivity across countries. Islam's work shows that once one accounts for differences in the level of productivity, the Solow model captures well the endogenous accumulation of physical capital, without any need to account for the accumulation of human capital. Islam suggests human capital's contribution to changes in growth may not be as evident because it changes so slowly: While physical capital may completely adjust to changes in tax policy in a matter of decades, human capital may require a century to respond to changes in educational policy!

Despite this evidence, Jorgenson continues, the proposition that private investment in physical and human capital is a more important source of growth than productivity remains as controversial today as it was in the early 1970s. Jorgenson believes that he has largely resolved this issue, however, with a perfectly competitive, constant-returns-to-scale neoclassical model that employs constant-quality indexes of both labor and capital input and investment goods output. The complete econometric model developed over many years by Jorgenson and his colleagues attributes fully 83 percent of growth to the endogenous changes in capital and labor inputs, with the remaining 17 percent accounted for by technological change and fertility rates. This finding essentially reverses the attribution of growth from that of Solow who found that only 12.5 percent of growth in per capita output could be attributed to capital accumulation (he did not consider human capital).

Discussant **Susanto Basu** assesses the success of the Jorgenson (and coauthors) research program according to its ability to explain three "fundamental questions of growth theory": (1) Why does per capita income increase over time? (2) Why are some countries rich and others poor? (3) Why has economic growth slowed down in developed countries?

With regard to the first question, Basu points out that Jorgenson treats technology as knowledge, which is a form of capital and behaves just like any other capital. The New Growth theory, by contrast, believes that the knowledge that propels technological advance differs from other capital in one crucial aspect: "Investors cannot fully internalize the benefits from accumulating knowledge." The presence of strong spillovers from private investment in knowledge can imply significant differences in the answers that Jorgensonian and New Growth theories give to the first question. The Jorgensonian rendering implies that in the very long run, no growth in per capita income can occur, since growth arises only from capital accumulation, and the marginal product of capital must diminish as capital accumulates. By contrast, the New Growth theory implies that the long-run growth rate of the economy will depend on the rate of accumulation of "knowledge" capital. Jones (1995) provides compelling evidence against the latter hypothesis for the United States and other advanced economies. Taking the inherent plausibility of knowledge spillovers together with Jones's evidence, Basu favors an intermediate position with modest spillovers, consistent with the Jones evidence and with the Jorgenson position.

The work of Islam (1995) highlights a deficiency in Jorgenson's approach with respect to the second question, namely that differences across countries in income per worker cannot be explained by differences in capital per worker, as required by the Jorgenson model. That is, countries' production functions cannot be the same. To explain income differences, we require another factor of production that varies across locations, perhaps a factor that involves differences in the diffusion of technology or the degree of infrastructure in place, and thus drives a wedge between technological change and productivity.

Could this wedge also explain the observed slowdown (since the early 1970s) in productivity in advanced countries? Basu suggests that it may. Using the methods of Basu and Fernald (1995), he presents estimates showing that only a small portion of the slowdown in productivity growth can be attributed to a reduction in the growth rate of technology. Basu suggests that changes in the allocation of inputs across sectors may account for the bulk of the productivity slowdown. He concludes by agreeing that Jorgenson's paper documents the explanatory power of the neoclassical model augmented by careful measurement. He believes, however, that the model will need to be amended to allow for some spillover effects.

Discussant **Gene Grossman** focuses on four key questions about the role of technology in growth theory. First, "Is technological progress needed to sustain growth?" Grossman notes that, technically, our economy could grow indefinitely without technological enhancements if we continue to invest in physical and human capital *and* if the returns to doing so always remain above a minimum level. However, he suggests that long-run growth with such static technology is implausible. In the presence of factors in fixed supply, such as land and fuels, capital must eventually experience severely diminishing returns. Would the world economy have evolved as it has over the past 200 years in the absence of all the innovations introduced in that period—without steam engines, electricity, or semiconductors? Adding more and more shovels and horses would not have allowed us to reach today's level of output. A role for technology in long-run growth seems mandatory.

A second question is whether innovation represents the product of intentional activity and is thus "endogenous" to the economy, or not. Grossman suggests that innovation is endogenous; the firms that spend in excess of $100 billion on R&D must be doing so for a reason. He also cites the evidence in Baumol that innovations vary across history in response to variation in incentives facing innovators.

Third, Grossman asks whether "formal" R&D is responsible for the bulk of technological progress. The evidence presented by Jones (1995) suggests not: The long-run surge in R&D activity in the postwar period has not been accompanied by equal surges in the growth of per capita output, and the decline in productivity since 1973 does not seem to be explained by declining R&D (Griliches 1988). Perhaps this mismatch of R&D and output growth reflects a focus on the use of "formal" R&D, which may not measure efforts to improve manufacturing processes or organizational structures, or, more generally, to innovate at the margin.

Finally, Grossman asks whether the market-determined level of R&D investment is socially optimal. The answer to this question depends upon the existence of knowledge "spillovers": Knowledge gained from one firm's investment makes research more productive for other firms, while the other firms need not compensate the originating firm for this knowledge. When spillovers exist, the *social* returns to investing, which include the returns to those who did not pay for the investment, exceed the *private* returns. Jorgenson is skeptical of the existence of such spillovers, but Grossman reads the bulk of the empirical evidence as pointing to social returns to R&D investment that are more than twice as large as private returns. Does the presence of excess social returns suggest an investment tax credit or subsidy to foster innovative activity? Not necessarily; as Mansfield (1986) points out, R&D tax credits often encourage firms to relabel existing activities as investment, rather than to undertake new research.

Grossman acknowledges the important contributions of the neoclassical framework, favored by Jorgenson, to growth theory. However, he points out limitations of the model that make it "not well suited for studying innovation": The neoclassical model assumes constant returns to scale and perfect competition. Investment in knowledge, on the other hand, requires large up-front fixed costs that imply *increasing* returns to scale, and pricing in excess of marginal costs to recover high fixed costs, in violation of the assumptions of perfect competition. Thus, Grossman feels, one must study innovation in a setting that allows for imperfect competition, even when this makes policy prescriptions more difficult.

Uncertainty and Technological Change

Nathan Rosenberg examines the relationship between uncertainty, technological change, and economic growth. Rosenberg's approach to the topic is, he admits, anecdotal; but he discusses many of the most important innovations of this century, demonstrating the influence of uncertainty for technologies that have had tremendous economic impact.

Many of Rosenberg's primary conclusions are exemplified in his study of the laser. The laser currently has dozens of applications, from producing CDs to enabling delicate eye surgery, from an essential instrument in chemical research to the rapid carrier of data, voice, and optical information across telecommunications lines. And yet the initial developers of the laser at Bell Labs not only could not foresee these applications, but did not think the invention worthy of a patent application, since "such an invention had no possible relevance to the telephone industry." This lack of foresight was not a malady unique to the telecommunications industry or to potential users of lasers; the same inability to predict the general usefulness of an invention, let alone its particular uses, extends to the developers of the telephone, the computer, the transistor, the jet engine, and the radio.

What categories of uncertainty make it so difficult to foresee the usefulness of innovations? Rosenberg catalogues several. First, new technologies arrive on the scene with characteristics that do not immediately or obviously lend themselves to application. For example, new techniques for visualization in medicine, such as CAT scanners and magnetic resonance imagers (MRIs), were developed before it was known how to interpret their output in a clinically useful fashion. Significant additional research was required to render the innovation not only technically feasible but also usable by doctors and technicians in making diagnoses.

A second class of uncertainty arises when the success of invention A depends on improvements in complementary invention B, which may not exist at the time invention A is introduced. Take, for example, the use of lasers in communications. Only upon the development of fiber optics, and upon understanding how laser light could be transmitted through fiber optic cable, did lasers become a viable communications medium. When the success of the innovation depends upon a system of complementary innovations, as may be the case with computer technology, the length of the gestation period from inception to a full menu of uses may be decades.

A third class of uncertainty arises because many inventions were designed to solve very specific problems. For example, British engineers invented the steam engine in the eighteenth century to pump out flooded mines. The possibility that such an engine could be used in entirely different industries, for transportation or power generation for manufacturing, became evident only after many decades, during which time a sequence of improvements were made to the initial invention.

Finally, Rosenberg identifies uncertainty about the marketability of an invention. As he puts it, inventions need "to pass an economic test, not just a technological one." When Marconi invented the radio, he did not possess David Sarnoff's vision of a new medium "to transmit news, music, and other forms of entertainment and information into every household in the country." Without someone to anticipate and champion the commercial possibilities of the technology, the radio might have gone the way of the buggy whip.

In concluding, Rosenberg draws out the policy implications of the almost overwhelming uncertainty involved in technological innovation. First, he suggests that the increased emphasis on the "relevance" of research to social and economic needs is misplaced; we cannot know which research or development will turn out to be relevant, or relevant to what! For the same reasons, the government should not attempt to support a single technological approach to a problem, or one narrow area of research. These caveats do not necessarily apply to the private sector, however. In the face of uncertainty, Rosenberg asserts, the market will of its own accord encourage individual firms to pursue a wide array of research strategies, which, given uncertainty, is more likely to produce a useful innovation.

Joel Mokyr is largely sympathetic to Rosenberg's characterization of the uncertainty (or perhaps ignorance) facing decision-makers, but he suggests a modest reinterpretation. First, Mokyr posits two levels of uncertainty in technological change, the firm's *micro*-uncertainty, and the economy's *macro*-uncertainty. The former comprises a host of firm-level questions: Can this particular technical problem be solved? Can this firm solve it? Will we arrive at the answer first? Will it sell, or sell profitably? At the macro level, uncertainty involves which technological regime will dominate: nuclear or fossil fuels? Both levels of uncertainty figure prominently in the decisions of potential innovators.

Mokyr poses an analogy between evolutionary biology and technological innovation. The analogy holds in two regards. First, innovations, like mutations, occur at least somewhat randomly, and thus we do not know in advance what the future supply of innovations will look like. The degree of randomness likely differs between biology and technology, as the latter presumably attempts to respond to economic need. However, Mokyr and Rosenberg agree that while the correlation between need and mutation "may not be zero, it is not very high either."

Second, we do not well understand the "laws" that determine whether a particular mutation will be *selected* or not, in the biological case by natural selection, and in the case of technology by the market. Success in many instances depends on luck; Mokyr points out that 70 percent of all new products that make it to the distribution stage disappear again within 12 months. This high mortality rate underscores the poverty of knowledge, even among the innovators themselves, about the laws that determine which innovations will be successful.

Mokyr adds a third "evolutionary" process that is germane to understanding the uncertainty in innovation: the evolution of economic institutions. As Douglass North (1990) has emphasized, institutions evolve in a way that is no more predictable than the evolution of science and technology.

But the situation is even more complex, as the sources not only evolve but *co*evolve. Many institutions—free labor markets, enforced property rights—are good for technological development, whereas others—uncertain property rights, totalitarian government—clearly are not. Modern innovators need to know how the institutional climate will be when they bring their product to market. Will the FDA approve it? Will I get sued? Will it pass environmental restrictions? Only as institutions friendly to innovation evolve with technology will technology succeed.

Mokyr concludes with reference to a final biological/technological debate, between "adaptationists" and "anti-adaptationists." Do technology and living species adapt so that we see only efficient technological and biological outcomes, or do important examples exist of innovations (mutations) that are clearly suboptimal and persistent? Is the dominance of the Qwerty keyboard a result of inefficient lock-in and path-dependence, or do we not properly understand its inherent efficiency? Mokyr declines to take a firm stance on this issue, but notes a difference between the biological and technology versions of the debate. The biological adaptation debate involves a more constrained evolutionary process: A species can adapt or become extinct. Technology is somewhat less constrained; societies can, at least in principle, adopt a completely different technology very rapidly, albeit at significant private and social cost. Does the private benefit to changing technologies cover the social costs of not changing? If not, another role for the government may be to spur such changes when private benefits fall short of total social benefits.

Luc Soete cautions against drawing broad conclusions from the anecdotal evidence presented by Rosenberg. The innovations chosen by Rosenberg may have sparked the interest of historians precisely because they had such unanticipated success; if so, they may not be truly representative. Soete also suggests that sectors vary greatly in the type of uncertainty facing their research efforts. A drug firm that pursues hundreds of leads on a trial and error basis faces a different kind and magnitude of uncertainty from a chip manufacturer that is developing the next generation that will double processing speed.

Soete questions whether omnipresent uncertainty could explain the productivity slowdown. Do the productivity gains that we expect from, for example, information processing technologies, seem to lag their invention because of the difficulties in identifying their most efficient uses? "You ain't seen nothin' yet" is the optimistic buzz-phrase of this explanation.

Soete proposes two other equally plausible explanations of the "missing productivity." The first is the difficulty inherent in measuring the output of information goods and services. As suggested by Nakamura (1995), the failure to properly capture the consumer surplus generated by the vast array of new electronic and communications products recently made available will likely underestimate output growth, perhaps by enough to account for the missing productivity. The second explanation centers on the possibility that the short-term disinflationary monetary policies of the 1980s, which significantly increased real long-term interest rates, may have turned businesses' research focus to short-term R&D with immediate payoffs, at the expense of longer-term, more uncertain research.

Cross-Country Variations in National Economic Growth Rates: The Role of "Technology"

J. Bradford De Long's paper attempts to explain two striking observations about the cross-country distributions of living standards and growth. The first is that the cross-country disparity of per capita real incomes has increased markedly over the past two centuries. The second is that the growth rates of real income in individual countries seem to be converging to the pace that is consistent with their rates of investment and population growth (as documented in the work of Ball, Mankiw, and Romer 1988).

Broadly construed, De Long's explanation works as follows. He notes that the countries that were relatively poor 200 years ago are relatively poor today, and those that were relatively rich 200 years ago are relatively rich today, and that the gap between the rich and the poor is increasing. According to the neoclassical model, if each country had started with somewhat different endowments of labor, capital, and materials but had access to the same technology, then over long spans of time, all countries would approach the same level of real per capita income. The long-run divergence of incomes argues against this simple case. If, however, the rich countries enjoy *amplified* effects of technology improvements on standards of living, while poor countries do not, then we will not observe even a gradual convergence of living standards.

De Long's paper identifies two novel sources of income divergence, each of which rests on a magnified long-run effect of productivity on real per capita income for richer countries. The first source is the strong endogeneity of population growth with respect to productivity and income. Countries with high productivity and thus high real incomes tend to have lower population growth rates. De Long shows that, for the United States, each tripling in real per capita GDP is associated with a 1 percentage point fall in the annual rate of population growth. De Long suggests several explanations for this pattern. More prosperous countries are often more educated countries, and better-educated women demand better birth control; in poor countries, the average number of years of schooling is low, and children are more valuable to production there because they can be put to work at an earlier age. In other words, children in poor countries are "investment goods" rather than "consumption goods," as they are in rich countries. Other things equal, then, a country that experiences rapid growth through increasing productivity will experience lower population growth that will, in turn, raise income per capita.

The second magnification effect arises from the endogeneity of the relative price of capital. Prosperous countries tend to benefit from a low relative price for investment goods. Most wealthy countries have achieved their prosperity largely through attaining high levels of manufacturing productivity. This achievement implies a relatively low price for manufactured goods, including the investment equipment that firms use to produce more goods. In support of the negative correlation between prosperity and the price of capital, De Long notes that the real purchasing power of domestic currency in foreign markets can be as much as eight times higher in rich countries than in poor countries. The disparity in real purchasing power directly reflects the difference between the relative price of easily traded goods, such as physical capital, in richer and poorer countries. This negative correlation between prosperity and the price of capital also magnifies the effects on real incomes of changes in productivity: As productivity and real incomes rise, investment goods become cheaper, and the economy can afford more investment goods for a given pool of savings, thus affording further increases in productivity.

De Long shows that the combined effect of these productivity magnifiers is substantial. Including them implies that the estimated effect of a productivity increase on the steady-state level of output is orders of magnitude larger than simple growth accounting would suggest. These important endogeneities between income, population growth, and physical investment could go a long way toward explaining the extreme divergence in national incomes that we have observed over the past two centuries.

Thus, De Long concludes that technology, broadly defined as differences in productivity, explains much of the disparity in standards of living across countries. He notes, however, that technology, narrowly defined as the possession of the most modern machinery and manufacturing processes by a particular country, explains relatively little of the differences in per capita incomes across countries. He cites work by Clark (1987) that shows remarkable differences in output per hour in cotton textiles across countries in the early twentieth century, even though many of these countries used exactly the same textile machinery. The McKinsey Global Institute's study (1993) of cross-country productivity differences reveals similar puzzles: Japan appears to be 47 percent more productive than the United States in steel manufacture, but 67 percent less productive in food processing. It seems unlikely that Japan is adept at using and refining the best manufacturing procedures for steel manufacture, yet is completely inept at "learning and developing technologies for making frozen fish."

Reacting to De Long's observation concerning the link between income and population growth, **Jeffrey Frankel** points out that "a prime motive in poor countries for having many children is that they provide the only form of insurance against destitution in old age." As a country develops, its financial institutions develop with it, and the increased accessibility of savings instruments can substitute for a high ratio of children to working-age population as a savings plan.

Frankel also observes that De Long's hypothesis about the endogeneity of both population growth and the price of investment goods suggests a timing test: Under De Long's interpretation, one ought to see significant decreases in population growth or increases in investment rates *following* surges in real growth. Frankel finds little evidence in the data for East Asian countries that declines in population growth are more likely to follow peak growth rates than to precede them. Investment rates follow peak growth rates in some cases, perhaps confirming De Long's hypothesis. However, the data also show large increases in investment that predate the peak in growth rates and could, thus, be considered the proximate cause of subsequent growth, contrary to De Long's interpretation.

Frankel ends by noting De Long's omission of a critical determinant of differences in growth across countries: openness to trade and investment. A large body of empirical work finds openness to be an important contributor to growth, even accounting for differences in factor accumulation. The economies that have converged are those that are open, whether across the OECD, across Europe, or within the United States. The reason, according to Frankel, is that "openness is how countries absorb the best technology from the leaders," whether we construe technology narrowly, as in the most up-to-date machinery and equipment, or more broadly, to include managerial and organizational techniques. In addition, openness to trade is part of a self-reinforcing pattern of growth: Countries that open their boundaries to trade grow more, but countries that have grown also tend to lower tariffs and promote trade.

Adam Jaffe presents cross-country evidence supporting the effect of income on population growth. Real per capita income and population growth exhibit a strong negative correlation, with an increase in per capita income from $1,000 to $10,000 associated with a decline in population growth from 2.5 percent per year to 1.5 percent. Of course, the link between income and population growth is partly mechanical: As population grows, holding income constant, per capita income must fall. But Jaffe shows that the strength of the correlation could not arise exclusively from this mechanical relationship. Suppose two countries begin with the same per capita income, but the population of one grows at 1.5 percent while the other grows at 2.5 percent. The low-population-growth country will reach an income 10 times the rapid-population-growth country only after 156 years! It is plausible, therefore, that much of the cross-section variation in income and population growth rates arises because high income causes low population growth, and not vice versa.

Jaffe suggests that the negative relationship between real income and population growth is not continuous. The correlation falls substantially for incomes above the median, and vanishes for countries with per capita incomes above $10,000. Thus, the returns (measured in lower population growth) to higher income appear to cease above this threshold income level. This observation alters De Long's story somewhat. Once the one-time demographic threshold is crossed, no further population growth effect would occur for the rich country.

Jaffe also clarifies the explanation for the observed correlation between income and the price of investment goods. Productivity improvements must (by definition) make goods and services cheaper. Because most of the productivity enhancements of the past century have been concentrated in manufactured goods, the real price of manufactured goods has fallen faster than the real price of services. As investment is likely to draw more heavily on manufactured goods than on services, the relative price of investment goods will also fall as productivity rises. The importance of this observation is that the apparent feedback between income and the price of investment goods can arise from productivity increases in an autarkic country, and thus does not depend upon foreign trade. The correlation between the real purchasing power of domestic currency and growth simply reflects underlying differences in productivity improvements across countries.

Finally, Jaffe questions the usefulness of a debate over which inputs to production should be labeled "technology." Echoing comments made by a number of participants during the conference, Jaffe finds it more useful to expand the list of potential explanations of differences in growth across regions and sectors. He suggests that a deeper understanding of the importance of hardware, software, human capital, ideas, and institutional and market factors in production may help us better explain differences in productivity and growth.

Address: Job Insecurity and Technology

Alan Greenspan's address focuses on human reactions to the structural changes caused by modern computer and telecommunications technologies. Pointing to the paradoxical pervasiveness of insecurity and malaise in a period of extended economic growth, restrained inflation, and a comparatively low layoff rate, he examines the origins of this anxiety and suggests ways of alleviating it.

He sees modern societies as having evolved from a time when the creation of economic value depended on physical brawn and physical product to the present when ideas are the critical input. This accelerating trend has had two important consequences: It has played a major role in changing the distribution of income in this country, and it has created a sense of foreboding in a large part of the work force.

Expanding on the first outcome, Greenspan explains that as ideas have become critical to the creation of economic value, education and intellectual skill have become increasingly important determinants of earned income. Although the supply of college graduates rose with demand in the 1960s and 1970s, by the 1980s the demand for skilled workers was apparently outstripping supply. The seeming result was a rise in the compensation of college graduates relative to that of less-educated individuals. Because the growth in real incomes slowed markedly in the mid 1970s—reflecting a similar (and not fully explicable) slowdown in productivity growth—widening income disparity has meant that parts of the work force have experienced stagnant or falling real incomes and understandably feel rooted to a treadmill.

Greenspan suspects that an even larger share of the work force is suffering from the job insecurity caused by rapid technological change. This group, composed of relatively skilled, experienced, and well-paid individuals who interact closely with our high-tech capital stock, are acutely aware of the speed at which this stock is being radically transformed. As a consequence, they fear that their own job skills may suddenly become obsolete. Greenspan suggests that these fears have led to an extraordinary period of labor peace, with a preference for job security over wage hikes, lengthening labor contracts, and unusually subdued strike activity.

Given widespread recognition of the growing income disparity, labor's acquiescence is somewhat surprising. Still, the relative economic welfare of low-income workers may not have deteriorated as much as the rising disparities in the distribution of income and wealth suggest. For example, recent work by Johnson and Shipp (1996) finds that the rise in consumption inequality since 1981 is only three-quarters as large as the rise in income inequality. Since purchases of consumer durables provide services throughout their useful lifetimes and are more akin to investments, the distribution of consumer durables deserves special attention.

Since 1982, household ownership of consumer durables has grown at an annual average rate of 3.3 percent a year, a slightly faster rate than in the 1960s and 1970s. Moreover, according to data provided by Stephanie Shipp and her colleagues at the Bureau of Labor Statistics, while ownership of consumer durables clearly rises with income, the distribution of ownership rates across income groups for cars and many appliances actually became more equal between 1980 and 1994. By exception, the disparity in ownership rates for personal computers remains large—unfortunately, given that knowledge of computers is linked to economic success.

Stressing that economic security depends on much more than owning selected consumer durables, Greenspan argues that the solution to the malaise created by rapid technological change involves finding ways to enhance skills. Since education has clearly become a lifetime activity, it is fortunate that many companies are beginning to see that human capital development is crucially important to improving profitability and shareholder value. He hopes that this approach will also help to reduce income disparities.

While the twenty-first century is likely to remain just as fast-paced as the recent past, Greenspan concludes, individuals currently entering the work force are used to rapid change and many six-year-olds are computer literate. Thus, as in previous periods of great structural change, the current frictions and uncertainties will diminish as people learn to adapt.

Microeconomic Policy and Technological Change

Reviewing the impact of public policies towards R&D spending, patents, and competition on innovation, **Edwin Mansfield** argues that government has a major influence on the rate of technological change in major industries. He points out that the federal government finances about 35 percent of all U.S. R&D investment and 60 percent of the R&D performed by colleges and universities. He provides two rationales for these expenditures. First, where government is the primary purchaser of public goods, like national defense or space exploration, the government clearly bears primary responsibility for promoting related technological change. In addition, much federal R&D is directed towards basic research because market failures or spillovers could cause private sector investment to fall short of socially optimal levels.

However, it is not self-evident that R&D spending is actually suboptimal. In many oligopolistic markets, product improvement is a major form of competition. As a result, R&D spending might actually exceed socially desirable levels in such industries. In addition, the government currently subsidizes R&D activities through the R&D tax credit and various grant programs. Thus, the government may already have offset any tendency for the economy to underinvest in R&D.

To address this issue, Mansfield reviews empirical estimates of the social rate of return from innovation, a body of work to which he has made major contributions. He starts by showing that the social benefits from an innovation equal the sum of the gains to consumers from the resulting decline in prices and society's resource saving (alternatively, the innovator's profit). Arguing that a high social rate of return signals a productive investment, Mansfield reports that empirical studies consistently find the median social rate of return from innovation to be substantial (the lowest median cited was 56 percent), even when private returns were low or negative.

The gap between the social and private rates of return from innovation provides an important rationale for government support of civilian technology. But, while a remarkable number of independent studies find the gap between marginal social and private rates of return to be sizable, many economists suspect that federal intervention could do more harm than good. Accordingly, Mansfield offers guidelines for public R&D support programs. First, given the huge uncertainty surrounding R&D outcomes, government incentives should remain modest, encourage parallel approaches, and provide information for appraising the desirability of further support. Such programs should not aid declining industries or late-stage development work. Recommending a pluralistic, decentralized approach, Mansfield also suggests that potential users of new technologies play a role in project selection so that public R&D efforts reflect market realities.

Mansfield's paper then reviews the pros and cons of another important instrument of national technology policy, the patent system. Some supporters argue that patent protection provides necessary incentives for innovation and development activities by slowing the introduction of relatively low-risk, low-cost copycat products. Other proponents assert that the patent filing process actually speeds the disclosure and dissemination of new technologies. Critics complain that the patent system creates usually weak but sometimes self-sustaining monopolies that slow the spread of new information. Still others conclude that patents have minimal importance, especially for large corporations; firms keep secret what inventions they can, they say, and patent those they cannot.

Turning again to empirical results, Mansfield reports that while patent protection does not make entry impossible or even unlikely, it does raise the cost of imitation. According to one study by Mansfield, Schwartz, and Wagner (1981), patenting raised the median imitation cost by 11 percent—30 percent in ethical drugs and 7 percent in electronics. Despite widespread skepticism about the value of the patent system, Mansfield acknowledges that few economists would recommend abolishing it, given our limited understanding of its impact.

Mansfield's paper ends with a discussion of the effects of market structure and antitrust policy. He concludes that while market entrants often play an important role in promoting technological change, some R&D activities exhibit economies of scale. Since a complementary mix of firm sizes appears to benefit technological change, public policy should aim to eliminate unnecessary barriers to entry and discourage industrial concentration.

The theme of **Samuel Kortum**'s comments is that the effectiveness of government technology policy depends crucially on the responsiveness of technological change to research effort, and that the evidence about the actual impact of research activity on innovation is weak. Although a vast literature has uncovered a systematic relationship between growth of total factor productivity and research effort (R&D/sales), Kortum points out that these studies provide no evidence concerning the direction of causality in this relationship.

Kortum raises the provocative possibility that technological change may be largely impervious to government incentives—if, for instance, innovation is an exogenous process more dependent on the chance arrival of technological opportunities than on incentives to exploit them—and sets out to show that this idea is not so easy to disprove. To do so, he develops a model in which R&D spending is the means by which firms compete for patent rights to innovations that arise within the economy regardless of the level of research activity. The larger a firm's share of industry spending on R&D, the greater is the probability that it will win patent rights valued at the industry's cost savings from the innovation. If the above model describes the real world, a cross-industry estimation of the impact of R&D effort on total factor productivity will reflect the fact that R&D effort depends on the value of exogenous innovation.

In Kortum's model with exogenous technical change, the private rate of return to R&D is the interest rate, but the social rate of return is −100 percent since the marginal expenditure has no benefit for society. Even careful economists, like Mansfield, who sum all research costs for losing as well as winning firms in calculating the social rate of return on R&D, are likely to find huge social payoffs—erroneously if innovation is actually exogenous. Although Mansfield and his coauthors state that social benefits should be measured only between the date when the innovation occurred and the date when it would have appeared if the innovator had done nothing, Kortum questions the validity of survey work based on hypothetical questions about the timing of competitors' innovations.

To provide additional evidence as to whether innovation is endogenous or exogenous, Kortum recommends careful analysis of the impact of a specific policy change, like the increased patent protection stemming from the 1983 creation of a single appellate court for patent cases. If technological change is actually exogenous, then such a policy shift should have no impact on productivity. By contrast, evidence that the policy action raised productivity would be highly suggestive of endogenous technological change.

Joshua Lerner focuses his comments on Mansfield's policy prescriptions. In particular, he asks whether technology policy should recognize that small firms generate a disproportionately large share of major innovations, since, as Mansfield and others have pointed out, many studies find that start-ups play a big role in applying radical technologies. Although key innovations are usually developed with federal funds at universities or research labs, small firms are often the first to act upon the commercial possibilities. As important examples of this phenomenon, Lerner cites the development of biotechnologies and the Internet. Given the uncertain path of technical developments and the critical role often played by previously unknown firms, Lerner is skeptical of Mansfield's stress on a "proper coupling between technology and the market" and his prescription that federal R&D be directed with the advice of potential users.

Lerner next addresses issues raised by the patent system, particularly the impact of the single court of appeals for patent cases established in 1983. Lerner argues that the new court has produced more pro-patent rulings than the previous system—with the result that large and small firms are putting more effort into seeking new and defending old patent protection. Viewed broadly, Lerner contends, the consequent growth in patent litigation has created a substantial "innovation tax" that falls particularly hard on small firms. In a recent research effort Lerner (1985) has found that patent litigation begun in 1991 will lead to total legal expenditures amounting to more than one-quarter of the private dollars spent on basic research; the indirect costs of this litigation are also substantial. Survey results suggest that these costs are a more important deterrent to development efforts for small firms than for large firms. Accordingly, reforms intended to protect and spur innovation have actually discouraged entry. Lerner is concerned that efforts to make federal research commercially relevant could have the same effect.

Technology in U.S. Manufacturing: The Geographic Dimension

Continuing with a micro perspective, **Jane Sneddon Little** and **Robert K. Triest** explore the process by which advanced technology enters general use. Using relatively new data from the U.S. Census Bureau's Surveys of Manufacturing Technology (SMTs) for 1988 and 1993 (U.S. Bureau of the Census 1989 and 1994), their paper examines variations in the adoption of 17 advanced technologies across the nation and within individual states. The authors consider a variety of plant and locational characteristics that might raise the probability of technology use, but they are particularly interested in whether proximity to firms already using advanced technologies fosters adoption. Proximity to early users might affect adoption decisions by reducing the perceived risk and actual cost of investing in this new equipment.

Little and Triest estimate a set of econometric models that control for the effects of plant, firm, and locational characteristics. As measures of technology diffusion, the authors examine the change in the number of advanced technologies used by SMT establishments between 1988 and 1993, the number of technologies used in 1993, and the probability of adopting a particular technology by specified dates covered by the SMT survey. In each case, the authors first control only for proximity to other users of advanced technologies. They then add in a set of plant and firm characteristics, such as size and industry. As a final step, they include a set of locational characteristics, like educational attainment of the work force, in the group of explanatory variables. In all three estimations, proximity to early users almost always has an economically and statistically significant positive effect on technology adoption, not only when proximity is the only explanatory variable but also when plant characteristics are taken into account. While introducing locational characteristics always reduces the coefficient on proximity, these coefficients still remain positive and statistically significant in the equations for the number of technologies used in 1993 and for the change in number of technologies used. By contrast, for the models estimating the probability of adopting specific technologies over a span of years, the proximity variable generally loses its significance when the geographic variables are added.

Little and Triest conclude that proximity to other users of advanced technologies is associated with higher rates of adoption, even when industry and other plant characteristics are controlled. They find this result noteworthy since, with its well-developed communications networks and national markets for capital goods and skilled workers, the

United States might be expected to approach the limiting case of immediate, costless diffusion of technology. Human capital appears to be an important part of the proximity effect, they speculate, because, among the locational variables, access to a work force with a high school education or some technical training is associated with a higher rate of technology adoption. Some of the remaining proximity effect may reflect the impact of social interactions in spreading technical information.

Although the authors were not able to separate proximity/spillover effects from the impact of educational attainment/university R&D to their satisfaction, they believe that the evidence of uneven technology diffusion warrants further research. Because technology adoption is extremely expensive for individual firms and the nation, gaining a better understanding of this process remains an important goal.

John Haltiwanger's comments on the Little-Triest paper center on his concerns about data and measurement issues and about the appropriate interpretation of their results. Citing recent research by Dunne and Troske (1995), Haltiwanger points out that the answers to the retrospective questions in the 1993 SMT on the timing of technology adoption appear subject to substantial recall bias. Respondents systematically date adoption more recently than was actually the case. As a result, Haltiwanger suggests, the Little-Triest variable measuring the change in the number of technologies used may actually be a better measure of the number of technologies in use in 1993. Thus, although Little and Triest find some evidence of clustering, the timing problems raise questions about the direction of causality and the underlying source of this clustering.[1]

Dunne and Troske's work raises another important issue, Haltiwanger contends. Their 1995 study finds evidence of significant rates of de-adoption for specific technologies. For example, for the matched sample of plants responding to both the 1988 and 1993 SMTs, 39 percent of the establishments using local area networks in 1988 were not using them in 1993. This finding suggests additional measurement problems or the intriguing possibility that firms experiment with new technologies that they eventually decide not to use. If so, a region that is relatively slow to de-adopt should not be labeled "advanced," Haltiwanger suggests.

Haltiwanger then takes up a line of argument similar to that raised by Samuel Kortum: Does the adoption of advanced technologies actually affect outcomes we really care about—the growth of income or

[1] In response to Haltiwanger's comments concerning their use of retrospective data, Little and Triest reran their regressions using the subsample of firms responding to both the 1988 and 1993 SMTs. Relying on current rather than retrospective data on technology use did not change the flavor of their results. If anything, the change strengthens the impression that proximity affects technology adoption. See Little and Triest, footnote 43, in this volume, for details.

employment or productivity? While one might presume such a connection, work by Doms, Dunne, and Troske (1995) suggests that differences in technology use are not particularly meaningful. Although Doms, Dunne, and Troske find that advanced technology use has a significant positive effect on plant-level labor productivity, differences in technology adoption account for only 1 percent of the total variation in labor productivity across plants. Moreover, these authors find no statistically significant link between technology adoption and the growth in plant-level labor productivity. (Perhaps the failure of micro studies to find much connection between the adoption of new technologies and productivity levels or growth should not be so surprising, given our similar inability to find any productivity payoff to vast investments in new technologies at the macro level.)

Alluding to research stressing the dominance of idiosyncratic factors and the importance of the reallocation process steering resources from less to more productive plants, Haltiwanger suggests caution in interpreting empirical results concerning technology diffusion. Seemingly, the growth process is noisy and complex, and the required resource reallocation is time-consuming.

In commenting on the Little-Triest paper, **George Hatsopoulos** provides the perspective of his many years of experience in managing high-technology companies. He interprets proximity as representing local management culture or standard technological practice within a given area. In this context he finds that the authors' conclusions correspond with his own observations.

Hatsopoulos starts by emphasizing the relative importance of diffusion—compared with innovation—in determining a country's technological sophistication. Like John Haltiwanger, he also finds that intangibles like managerial and organizational skills, and labor-management relations, exert an extremely important influence on micro and macro productivity levels.

Turning to Little and Triest's empirical results concerning the impact of proximity on the probability of technology adoption, Hatsopoulos reports that this finding matches his observation that decisions about the use of specific technologies are determined by middle managers and foremen who, in turn, are heavily influenced by prevailing practice at neighboring plants. To illustrate this point, he cites the example of two plants, one in Manchester, England and one in Auburn, Massachusetts. Although the two were making identical products for the paper industry, labor productivity in Manchester was about half that in Auburn. The problem, it turned out, was that managers and workers in Manchester were extremely reluctant to import manufacturing and organizational technologies that headquarters had found useful in the United States but that were uncommon in Britain. Because these workers were very heavily influenced by local practice, Thermo-Electron had a very hard time trying to change their behavior.

Reacting to Little and Triest's finding that plant size has a significant positive impact on technology adoption while firm size does not, Hatsopoulos indicates that these relationships again appear intuitively plausible to him since plant scale must be considered in making technology decisions while access to capital, a firm-level characteristic, has only an indirect impact on local technology choices. Similarly, Hatsopoulos reports that he is not particularly puzzled by the result that access to a work force with a high school education has a greater impact on the probability of technology adoption than does access to workers with a college education. Because he finds the importance of foremen and other middle managers to be of overriding importance in the technology decision, Hatsopoulos finds this result matches his expectations.

Macro Policy, Innovation, and Long-Term Growth: A Panel Discussion

Martin Baily begins by dissecting potential GDP growth, estimated to be 2.3 percent per year, into its major components: labor inputs, which have been rising about 1.1 percent annually; and labor productivity, which has shown trend growth of 1.1 percent per year since 1973. As Baily points out, while trend labor productivity has fallen from its 2.9 percent average in the 1960-73 period, the explicable part of productivity growth (the part due to capital intensity, education and experience, and R&D) has been remarkably constant at 1.1 percent since 1960. By contrast, the unexplained residual, the productivity "bonus" enjoyed between 1960 and 1973, has entirely disappeared; "We did not know where it came from then, and now we do not know where it has gone." In a related puzzle, the growing gap between the annual earnings of college and of high school graduates is widely attributed to a rising demand for technically skilled workers, but we see no signs of major technological breakthroughs in the productivity numbers. More formally, we see evidence of technological bias in the increased return to education but no evidence of technological change in measured productivity growth.

Turning to policy prescriptions, Baily concludes that current growth rates reflect supply rather than demand constraints and, thus, that the potential role for monetary policy in spurring growth is limited. By contrast, fiscal policy is important: During the 1980s the federal budget deficit was a primary cause of our low rates of saving and investment, which in turn contributed to the deceleration in capital intensity and productivity growth. Thus, reducing the federal deficit remains an important policy goal. A second area for policy action relates to education and training. Although the contribution to productivity growth made by education and experience has risen recently, that increase merely reflects the growing experience of the aging baby boom generation, and the rising return to education suggests that the demand for highly skilled workers continues to outstrip supply. Because Alan Krueger's work (1993) shows that computer skills in particular are linked to higher wages and, presumably, thus, to higher productivity, federal seed money for computer literacy programs might prove especially helpful. Finally, since studies by Edwin Mansfield and others suggest that the social return to private R&D is substantial, Baily concludes that tax incentives for R&D could play a positive role. Moreover, since private R&D appears correlated with prior federal R&D spending, Baily is concerned about congressional proposals to curtail the rate of public non-defense research.

Baily ends by speculating about the unexplained growth bonus enjoyed between 1960 and 1973. Much of that spurt in productivity growth may have resulted from a burst of innovation and a shift from craft to mass production that cannot be repeated. If so, we may simply have to adjust to a world with lower productivity growth and slower growth in average real wages—a world split into winners and losers. Such a world would require attention to policy dimensions such as the provision of safety nets, Baily submits. On the other hand, because measuring output and productivity is extremely difficult, particularly in areas like health care, or in retailing and financial services where convenience is important, output and productivity growth may actually be better than we think. Accordingly, Baily advocates investing in, not starving, our statistical agencies in order to get better data and better policies. Finally, maintaining open economies and deregulating domestic markets provide important incentives to adopting better technologies.

Ralph Gomory addressed his remarks to the impact of economic development in technically backward countries on welfare in the industrialized nations, a topic of great concern to many policymakers. As underdeveloped countries improve their technical capabilities, they become significant contributors to world output, but they also become more effective competitors to established industries in developed nations. What is the net impact on the national welfare of the technically advanced nations? To analyze this issue, Gomory offers a classical Ricardian model of international trade in which the relative efficiencies determining comparative advantage are allowed to vary, as in Gomory and Baumol (1995a).

Gomory sketches a two-country model—or rather a family of two-country models—that assumes single-input linear production functions, Cobb-Douglas utilities, and fixed labor supplies and demand parameters, as well as a fixed number of industries. In equilibrium, both countries actively participate in a given industry *only* if their unit labor costs in that industry are equal. The exercise then calculates, for all possible values of average labor productivity, the equilibrium outcome in terms of national utility and share of world income for each country.

The results suggest opportunities for inherent conflict between the two countries because, even though *world* output is greatest when both countries have similar productivities and split world production 50-50, the best outcome for each one *singly* occurs when it has a large share of world output and income; this point always represents a poor result for the other. As Gomory carefully points out, improvements in productivity in one country (which always increase that country's share of world income) sometimes enhance welfare in both countries; however, in other cases, unilateral improvements in productivity decrease the welfare of the other.

What conditions determine the outcome? Assuming, as in Gomory and Baumol, that efficiency rises in an active industry and decays in a less active industry, the model suggests a natural tendency for national shares of world output to remain close to their original values, while incomes expand as a result of improving efficiencies. However, if one country (generally the lower-wage country) succeeds by policy measures in "capturing" a growing share of world output in a given industry, its welfare improves. Whether or not welfare improves in the second (advanced) country depends on whether the depressing effect of the capture is or is not outweighed by improved efficiencies (via learning-by-doing, for instance) in all other industries. This result contrasts with Ricardo's original insight that trade based on comparative advantage determined by a specific pair of production functions always enhances well-being in both countries.

Abel Mateus's experiences with the Banco de Portugal and the World Bank permit him to examine the impact of macro policies on growth from the perspective of developing as well as developed economies. He suggests that technological progress is a primary determinant of growth in developed countries, whereas in developing countries most growth is due to the accumulation of physical and human capital that incorporates ideas transferred from advanced nations; thus, in these developing countries, outward orientation is complementary to the capital accumulation process.

Mateus points out that in small open economies "miraculous" growth is linked with rapid accumulation of human capital and use of that knowledge to operate physical capital to produce goods near the country's technological frontier. Shifting labor and capital to ever more advanced activities allows learning by doing and augments the accumulation of human capital. Export orientation is essential to such a growth strategy because this approach creates a gap between the mix of goods consumed domestically and the mix of goods produced and by necessity exported to larger, more demanding foreign markets. By contrast, Eastern Europe provides counterexamples of countries where the technology gap is sufficiently huge that trade promotes so much Schumpeterian (creative?) destruction that short-term welfare actually declines. Nevertheless, Mateus argues that these "industrialized" transitional nations must pursue the painful path of institutional change, industrial restructuring, and integration into the world trading system. Moreover, most developing countries, with smaller initial manufacturing sectors, do not face such conflicts; for them, the benefits of trade based on comparative advantage apply even in the short run. The policy implications stemming from Mateus's observations of small open economies include an emphasis on formal education, protection of property rights, and an export-oriented trade stance to promote competition and technology transfer.

Mateus then addresses the impact of free trade in goods and technologies on the developed countries, where these developments have been associated in the 1980s and 1990s with high unemployment rates and stagnant or declining real wages for unskilled workers. After a reminder that present levels of global integration are not unprecedented, he points to the drop in transportation costs and the increase in communication speeds as the truly new elements. Although he sees some evidence supporting Paul Krugman's (1981) hypothesis that these developments will improve the lot of peripheral regions at the expense of the core and Jagdish Bhagwati's finding that comparative advantage has become "kaleidoscopic," moving almost at random across developed countries, he finally concurs with Obstfeld (1994) that financial integration, with investment shifting from lower-return to higher-return projects, can yield substantial welfare gains throughout the world via its effect on output and consumption growth.

Because the profitability of innovation and diffusion depends in part on the macro environment, Mateus then turns to fiscal policy and suggests that a high and rising debt ratio is likely to lower the long-term rate of growth. He cites World Bank findings that a 1 percentage point increase in the government surplus as a percent of GDP raises per capita growth by 0.37 percent and the investment ratio by 0.24 percent. Other research suggests that debt ratios and budget deficits are positively associated with increases in long-term risk premia. Mateus concludes, thus, that the near doubling in gross public debt as a share of GDP between the 1970s and the mid 1990s has had a significant negative impact on European growth rates. Accordingly, Mateus recommends wider use of consumption-based tax systems and a significant cut in the size of the public sector, to be accomplished, in part, through better project and activity evaluations. In addition, Mateus warns, social security systems in most countries are unsustainable.

Mateus ends by reprising his major policy recommendations. First, the emphasis on economic stability, trade liberalization, market-oriented policies, and human capital accumulation long advocated by international organizations appears to be appropriate. Second, the potential for improving world welfare by technology diffusion and portfolio diversification is enormous. Finally, within the developed world, blaming globalization and "social dumping" for current labor market problems is misguided. In Europe, reducing high rates of unemployment requires improving labor market flexibility, while in North America dealing with the plight of unskilled workers awaits more adequate redistributive policies.

Robert Solow expressed relief that the panel was discussing whether macro (not monetary) policy can promote long-term growth; as phrased, the question implies that fiscal policy is available for the task—luckily, since monetary policy cannot possibly address the many goals often assigned to it. Solow then begins his policy recommendations by urging advocates—academics as well as politicians—to stop making inflated claims for their favorite policy tools. The flat tax, a cut in the capital gains tax, and various labor market reforms may or may not be good ideas, but their impact on growth is likely to range from negligible to small—at most. In particular, Solow chides, too many theorists have taken to fabricating powerful policy options by leaping from empirically established links between levels to assumed links between levels and growth rates. For example, while most would agree that the level of human capital affects the level of output, too many go on to assume that a high *level* of schooling will increase the *growth* of human capital, or that a high *level* of R&D will speed the *pace* of innovation. With these assumptions, tax policy can readily be shown to affect the permanent rate of economic growth since it is quite easy to design incentives for schooling or R&D. "But do we really know that an increase in schooling or R&D will generate more than a one-time shift in the level of output?" the self-described spoilsport asks.

This plea for circumspection limits the list of growth-promoting policies severely, Solow admits. Still, he considers certain commonplaces worth repeating. Given how little we know about the links between stocks and growth rates, any policy that raises potential output permanently should be described as contributing to growth—even if the long-term rate of growth remains unchanged. Just shifting the steady-state growth path upward, parallel to itself, is a major feat, he contends.

After warning that the trade-offs between growth and current living standards must be weighed, he emphasizes that anyone choosing growth must favor investment over consumption. Since a pro-saving policy need not be pro-investment (because additional saving may reduce a current account deficit rather than raise investment, say), Solow proposes combining improved incentives to save with policies that shift the composition of demand in favor of investment. Any fiscal stance, he reminds us, can be weighted in favor of investment, with tax-and-subsidy policy an obvious instrument. Similarly, while a given macro posture can be achieved with many combinations of monetary and fiscal ease and tightness, in general growth is efficiently served by mixing relatively tight fiscal policy, to promote national saving, with relatively easy monetary policy, to spur domestic investment.

Solow also endorses a macro strategy that guides total demand toward potential whenever a gap between the two appears—for many reasons, but not least because this policy is growth promoting. He notes in this context that actual demand tends to fall below potential somewhat more often than it exceeds it, and that prospects for weak and fluctuating demand discourage investment. While the impact of modest overheating (particularly on investment volumes) is less clear, he cites consensus views that price stability encourages the most productive allocation of capital. Solow ends by asking, tentatively, if the Fed could usefully conduct open market operations at all maturities, not just at the short end, in order to affect long rates which, presumably, are the most relevant for investment decisions.

Moderator **Richard Cooper** initiated the general discussion by remarking that over the last 50 years, the process of innovation has, for the first time, become institutionalized and by asking the panelists and conferees to consider the price of future growth in terms of current income. Would it have been moral to ask our grandparents to save more in order that we could be even better off, compared with them, than we already are? In response, Robert Solow replied that he would be less concerned about growth if we were better at income redistribution, but, since we find redistribution hard, increasing today's growth is one way to help today's poor children and today's poor countries. Baily and Mateus added that public and private myopia about looming retirement needs requires current policy action to spur saving. Other comments addressed the differential impact of environmental spending on measured productivity and the quality of life, and the need to explore the impact of the transitional costs of technological change on the growth process.

Conclusion

After two decades of research focusing on the source and stabilization of short-run economic fluctuations, the profession has recently returned to considering the determinants of long-run growth. This resurgence in interest arises for several reasons. Many developed economies have seen their average growth rate halved since the mid 1970s, and as yet we have no compelling explanation. Differences across countries in standards of living and in growth rates are large and not obviously shrinking, even as modern technology has been disseminated more widely and educational standards have risen. The welfare implications of these cross-time and cross-country income differences dwarf those that arise from business-cycle fluctuations.

One fruitful vein of research has striven to understand growth from within the neoclassical framework, attributing continued increases in income primarily to investments in physical and human capital. Dale Jorgenson's research constitutes probably the most carefully measured and estimated set of econometric models in the neoclassical tradition. His conclusion is that investment can account for the preponderance of growth. This assessment is important, as it provides a benchmark for the contribution of standard inputs to growth. And yet, as Susanto Basu points out, the neoclassical model ultimately cannot plausibly explain all of the differences in growth that we observe over time and across countries. It seems extremely unlikely that we could have achieved most of our high standard of living today simply by using more and more of the investment goods that were prevalent in the nineteenth century. We could not have arrived at our sophisticated communications-linked, information-processed, efficiently manufactured state simply by using more and more shovels, adding machines, and steam engines. And yet the available data do not reveal a clear relationship between the invention, development, or adoption of new technology and subsequent improvements in productivity or income. Where does this observation leave us? Participants in this conference generally agreed on a few tentative conclusions.

First, it may be helpful to understand the input to production that is neither human nor physical capital not simply as "technology," but as an aggregate of the state of technology, organizational and managerial ability, and "economic culture." These concepts are not easily measured, but given the inability of relatively well-measured constructs to explain the variation in productivity in disaggregated data, we must try to model and measure these intangibles better if we are to understand significant differences in growth and productivity over time and across countries.

Second, most conference participants agree that it is probably beyond our grasp to design policies that we can be confident will spur specific innovations, or even spur innovation generally. The difficulty arises largely from the tremendous amount of uncertainty that surrounds the process of innovation. Given the difficulty in knowing which innovations will succeed, when they will arise, and what complementary innovations they will require to become "useful," policymakers do not possess the foresight to tailor policies to foster specific innovations. Still, most participants agreed that the social returns to innovation exceed the private returns. Although the extent to which private returns spill over into non-appropriable social returns is not clear, most would say such spillovers are likely to be sizable. Thus, the government should play a limited role in promoting R&D.

Finally, two clear insights from our panelists merit special attention as pointers to future research. The first, highlighted by moderator Richard Cooper, is that we assume, as a matter of default, that a higher long-run growth rate is better. In doing so, we are implicitly choosing the multiple by which our descendants' welfare will exceed our own. At a 1 percent rate of annual productivity growth, our grandchildren will on average have 65 percent higher real incomes than we do; at a 2 percent rate they will have nearly triple our real incomes. But in order to attain these increases for our descendants, we must forgo some current consumption. Cooper poses the question: How much better off should our grandchildren be than we are, and at what cost? Robert Solow points out, in response to Cooper's question, that productivity-generated increases in the size of the economic pie may benefit the poor children of today *and* tomorrow. This question lies at the root of the discussions about productivity slowdowns and hoped-for improvements.

The second insight, articulated by Robert Solow, is a reminder that not all improvements in welfare must be measured as changes in the growth rate of the economy. One-time permanent improvements in the *level* of potential output are also valuable and probably much more attainable. The profession may do well to focus more of its attention on policies that could more reliably achieve these less spectacular improvements.

References

The references cited in the Overview can all be found in the papers and discussions, except for the following.

Krueger, Alan B. 1993. "How Computers Have Changed the Wage Structure: Evidence from Microdata, 1984–1989." *Quarterly Journal of Economics,* vol. 108, no. 1 (February), pp. 33–60.