Recent Productivity Growth: The Role of Information Technology and Other Innovations

Martin N. Baily

Senior Fellow
Institute for International Economics

This article is based on the keynote address delivered in November 2003 at the conference organized by the Federal Reserve Bank of San Francisco's Center for the Study of Innovation and Productivity on "Technology, Productivity, and Public Policy."

I am pleased to be talking about one of the truly salient issues for the economy today to such a distinguished audience. The acceleration in productivity growth since the mid-1990s has been an exceptionally important development for the U.S. economy. My remarks are going to run somewhat contrary to the view that information technology has been the overwhelming impetus to the acceleration in productivity growth in recent years. While technology and in particular information technology (IT) are obviously very important drivers of productivity, I want to make the case that perhaps they have been overemphasized in the literature. This might sound a little subversive in the context of this conference; but it really is not, since my comments are entirely consistent with the broader theme running through the sessions, that productivity is linked to innovation.

A considerably less controversial point now is that productivity growth has accelerated. Obviously, there was a period when we wondered if the increase in productivity growth starting in the mid-1990s was just a cyclical effect. Now that concern seems to have gone away, in part because measured productivity growth now is even faster than it was from 1995 to 2000. Whether that will survive data revisions and changes in measurement methodologies remains to be seen. However, unless there are drastic revisions, it certainly looks as if the higher rate of productivity growth is staying around for a while.[1]

Has the Contribution of IT Been Overstated?

One of the reasons that the contribution of IT to growth may have been overstated is that the main tool by which we judge such contributions is growth accounting. This is obviously a valuable tool, but it has some pitfalls; in addition, it requires careful interpretation as to what it actually says about the recent growth period.

In concept, growth accounting is straightforward—it is a framework for allocating the growth in output to the growth in capital (plants, equipment, and software), the growth in labor, and the growth in total factor productivity (TFP). In principle, growth in TFP stems from improvements in the ways capital and labor are combined in production processes. This might be thought of generally as changes in how workplaces are organized and managed.

In practice, however, we know that measured growth in TFP is affected by the errors in measuring capital and labor inputs. The potential mismeasurement of IT investment is one of the pitfalls of relying on growth accounting that I will come to in a bit. First, I'd like to discuss other reasons to be careful about the empirical results relating to the 1990s experience that use growth accounting measures, especially ones using aggregate data for a single country.

One reason for caution is that the growth accounting results from the 1990s reflect the fact that there was essentially a "one-observation" correlation in the aggregate data: IT investment and labor productivity accelerated at about the same time. This convinced many people that the two events were related as cause and effect—the productivity acceleration must have been driven by IT investment.

1. Fuller discussion of the issues discussed in this article can be found on the website www.iie.com.

However, before the 1990s productivity speedup, there was the 1970s productivity slowdown. Measured labor productivity growth went from an average of 2.8 percent during the period spanning the 1950s through the earlier 1970s, to an average of about 1.4 percent from 1973 to 1995. During the sessions at this conference, there were discussions about the host of reasons given for the slowdown. The one that was certainly the most popular when it was proffered in the 1970s was that the productivity growth slowdown was caused by the rise in oil prices. At that time there was a different one-observation correlation: productivity slowed down and energy prices went up at about the same time. Many people, me included for a while, believed that higher energy costs had caused slower productivity growth. It was not a crazy idea; indeed, there were good arguments, such as that rising energy prices would cause labor productivity growth to slow down because labor would be substituted for energy-using capital. There were also objections to that story, of course, but what really caused the collapse of the energy–productivity story is that oil prices came back down in the 1980s and there was little or no improvement in trend productivity growth. That experience should sound a warning that one-observation correlations can easily collapse when a second observation comes along.

Another reason for caution in linking IT investment to the post-1995 productivity acceleration is that the growth accounting framework did not do very well in tracking productivity movements prior to 1995. There was a fair amount of IT investment going on well before 1995. Indeed, that investment inspired the famous Solow paradox—Robert Solow quipped in 1987 that he could see computers everywhere except in the productivity data. Moreover, we now seem to be dealing with the reverse of the paradox: there was a severe slump in technology spending that began in 2000 and lasted into 2003, yet productivity growth has remained strong and may even have accelerated; from the first quarter of 2001, which is roughly the official cyclical peak, through the third quarter of 2003 productivity growth averaged over 4 percent a year. Later on I will talk about how investment in intangible capital may help, at least in part, to resolve this paradox. However, as with the oil price story, on the surface there appears to have been a breakdown in the one-observation correlation between the productivity growth trend and its apparent driver.

A second concern about taking the growth accounting results at face value is the possibility of reverse causality in the 1990s. To a degree, it may have been the economic boom that created the big wave of IT investment rather than the wave of IT investment that created the boom. Investing in IT became the thing for companies to do. Profits were strong and IT managers had considerable influence over investment decisions. Hardware and software vendors hyped the benefits of their products, and no CEO wanted his or her company to be left behind in the high-tech boom. Some of that investment did pay off in higher productivity, but there may have been a systematic tendency to overinvest. Expected returns were higher than actual returns over this period—hence, in part, setting the stage for the technology bust after 2000.

This latter concern, of course, is not directly related to any shortcomings of growth accounting, per se. Indeed, growth accounting is part of the basic tool kit of economics and is firmly established as a valid approach. So let me explain why I have reservations about taking its findings at face value. Here we get back to the issue of mismeasurement. The first question is whether or not to believe what the computer price deflators are telling us. Is it really true that quality-adjusted computer prices are falling rapidly, or are there problems in the methodology used for their estimation? There have been questions raised about the hedonic approach, but this does not seem to be a key issue quantitatively. The computer price deflators are built not only from hedonic regressions but also from matched models. The matched model approach follows the price over time of the same computer or compares prices in different periods for computers that are functionally very similar. In addition, the price indexes based on the matched model approach have fallen about as fast as the prices based on hedonic regressions. It really does seem to be the case that quality-adjusted computer prices have fallen extraordinarily fast.

Nevertheless, it still is fair to ask whether the price deflators accurately reflect the effective use-value of IT capital in the production process. Many of us have voiced skepticism based on anecdotal evidence. Robert Gordon (2000) has argued that the latest version of Word run on the latest model computer does not do that much more than the previous version and does not increase the computer user's productivity by very much. Going beyond anecdotal evidence, one of the McKinsey Global Institute (2002a) case studies examined what happened when banks invested in new generations of faster, more powerful computers. The McKinsey team found that most of the lower-skill and middle-level employees (tellers, for example) were doing pretty much the same things in the same way with the new computers as with the old computers. The banks reported little productivity gain in their retail operations from upgrading their computer hardware.

If this case study evidence is correct, however, it poses an important puzzle. If the new computers and software yielded little in the way of productivity benefits, why did the banks invest so much in them? Both the hedonic approach to computer price measurement and the matched model approach reflect what customers (businesses or individuals) are willing to pay to get the latest technology. If people are willing to spend their own money to buy new hardware, they must expect a gain from the purchase. Is there some systematic bias or failure of expectations that could explain this puzzle? Or are we missing a dimension of the decisionmaking?

Part of the answer may be that it is difficult to infer the use-value of investments when there are network effects, and computer technology is one area where these effects are important. The value to a business (or to a consumer) of a computer depends a great deal on the compatibility of that computer with other people's computers. One reason that banks put very high-powered computers on the desks of all of their employees is that they follow replacement cycles. After a period of time some computers wear out and, more importantly, some computers become obsolete in that they are not able to meet the high-end technology needs of some employees. To maintain internal compatibility, a company undertakes a comprehensive upgrade so that the speed of the convoy, if you like, is pulled along by the fastest ship rather than going at the speed of the slowest ship. The reason a low speed computer enters with a very low price when the Bureau of Labor Statistics collects its data is that many firms cannot choose to buy the less powerful computer since it would not be compatible with their other computers or with the software in use throughout the company. And the issue of compatibility may also extend outside a specific company to its suppliers or customers.

This argument is suggestive but not proven. Still, in the presence of strong network effects, it is very difficult to assess the benefit of IT investment to one part of a business enterprise without considering the enterprise-wide effects. In the case of IT investment by banks, for example, we need to know the impact on productivity if lower-skilled and middle-level employees were to continue to use less powerful computers while the rest of an organization upgraded its IT hardware and software.

Insights from the Use of Disaggregated Data

The severity of the pitfalls for empirical analysis I have mentioned are mitigated to some extent in studies that use disaggregated data. There are a number of excellent studies that use industry level data, but rather than going through a full literature review, let me mention work by Kevin Stiroh and by Dale Jorgensen, along with a paper that I wrote with Robert Lawrence.[2] These papers suggest that industries that had invested in IT hardware had indeed achieved faster productivity growth than other industries in the 1990s and had achieved greater productivity acceleration after 1995. Not all the industry evidence point the same way, however. In particular, an article by Jack Triplett and Barry Bosworth (forthcoming) finds that the industries that had surges in labor productivity growth were not necessarily the industries that had surges in IT investments. Rather, more of the post-1995 increase in labor productivity growth on an industry level seemed to come from TFP than from investment in IT. They were using the growth accounting framework at the industry level and, not surprisingly given the aggregate results, they did find that IT contributed to productivity in the 1990s as a whole. What they did not find at the industry level was that IT contributed especially to the post-1995 acceleration.

Going below the industry level, a study by Lucia Foster, John Haltiwanger, and Cornell Krizan (2002) finds that the exit of low-productivity establishments and the entry of high-productivity establishments generated most of the gains in retailing. That is a strong result, and it raises questions about the impact of IT. It suggests that investment in IT within existing retail establishments was not the main source of productivity increases in that industry. This is a notable result since retailing is a sector that has made one of the biggest contributions to overall productivity growth in recent years.

Key Drivers of Innovation: IT and Competitive Intensity

Detailed industry case studies also can shed light on the role of IT, especially how it is used in the business setting. I will use the case study evidence from the work carried out by the McKinsey Global Institute (2002a). Let me reiterate: I do not want to make the argument that IT investment was all a failure or that it contributed nothing to productivity growth. My comments are intended to shift the debate away from an uncritical acceptance of the idea that IT investment has been the dominant driver of faster productivity growth.

The analysis from the McKinsey case studies suggests that there is not one simple paradigm that can describe how IT affects business production processes and productivity. They looked at examples where IT investment did have a positive impact on productivity and found that the IT applications were used in very different ways across industries and even within subcategories of an industry. For example, in retailing, the mass-market supercenter, big box discounters, such as Wal-Mart and Costco, used IT to improve their supply chains and to allow the stores to manage their inventories, purchasing, etc., more effectively. Among spe-

2. See, for example, Stiroh (2002), Jorgensen, Ho, and Stiroh (2004), and Baily and Lawrence (2001).

cialty retailers, such as J. Crew and The Limited, IT allowed the stores to track changing customer tastes and match their own inventory mix to stock the right products for their customers (making sure the "hot" products are on the shelves). In both of these examples, IT is used to improve the supply chain, but it is serving a different role in each. For the discounters, its purpose was primarily to cut costs, allowing the stores to sell at lower prices. For the specialty stores, it allowed them to sell higher margin items with greater value-added.

Earlier I talked about banks and the apparent low returns they obtained from one part of their IT investment. Some of their investments clearly succeeded, however. Over the course of the 1990s, banking organizations saw a massive increase in the number of phone requests and inquiries from their customers (the number of phone inquiries rose from just over 1 million in 1994 to 2.3 million in 1998, a 24 percent annual rate of growth). The banks made major investments in computerized voice response units to handle the growth of calls. Customers would rather talk to a person than a machine, but without the greater use of IT, banks would not have been able to provide the same level of service profitably. There was, in fact, a large increase in the number of people employed to handle calls, as well (employment in this activity grew at 13 percent a year between 1994 and 1998). But it would have taken an even larger increase in personnel to handle all of the calls personally. With the IT investment, banks were able to expand their services and, at the same time, sustain overall labor productivity.

The design of chips is another example where the use of IT had a positive impact on productivity, but the way that was achieved is very different from the two previous examples. In developing a new chip, the constraint on performance improvement in the late 1990s was not merely the number of transistors that could be packed onto a chip. Over the long term, the ability to manufacture chips with greater technological capability (to keep Moore's law going) is obviously important. But at that time, the binding constraint was the ability of companies to design chips that took advantage of the capability that was available. The complexity of chip design had increased very rapidly, almost too rapidly for companies to deal with effectively. Consequently, it was developments in electronic design and testing tools that helped companies such as Intel improve productivity.

These are all examples of how IT investment paid off in higher productivity. There were a number of cases where the investment did not pay off. One example was already discussed, as banks upgraded the computers on the desks of tellers. Another example is investment in customer relations management software. As companies installed this software, they planned to collect reams of data on their customers that would allow them to customize their services. For example, customers would walk into a hotel, where employees would greet them personally, equipped with detailed information about what kind of rooms they liked and what services they might want. In some cases, the hardware and software were purchased but were never really put into operation because of data problems or other difficulties. In other cases the hotels were not able to generate higher revenues through using the systems—or at least they were not able to earn a positive return on their investment.

Another innovation, enterprise resource planning (ERP) software, has had mixed results. It has had some successes, as evidenced by the growth of SAP, one of the leading providers of this software. But many of the efforts involving ERP did not improve productivity. A key shortcoming is that ERP tends to lock in a certain way of doing things making innovation more difficult. Another problem with ERP is that the technology has had difficulty dealing with the degree of data incompatibility within many firms, so the promise of total information integration frequently was not realized.

The case studies, then, suggest that IT has paid off in some cases, but not in others. That is to be expected; after all, in any area, some projects will pay off and others will not. However, three important lessons emerge from the case studies. First, the examples where businesses were successful in their IT investment were ones where the companies had already identified a specific problem to be solved or opportunity to be exploited. They then looked at their business system to figure out how it could be changed to solve the problem or exploit the opportunity. Then, as part of that process, they figured out how IT investment could contribute to the success of the improved business system. Successful IT applications also typically occurred where outcomes could be quantified and monitored. The second lesson, as I stated earlier, is that when IT works, it can work very differently from industry to industry and even from firm to firm within an industry.

The third lesson is that, for many productivity-increasing innovations, IT investment was not the most important component. Take the case of the retail sector. Generally, companies can obtain labor productivity increases by building larger stores, such as the big box stores. They can also lower their costs by efficient national and global sourcing where, again, scale is an important benefit. Large purchasers are able to obtain lower prices from suppliers, thereby adding to their margins and value-added per employee. In retailing and many other industries, figuring out what consumers want and what they will pay for allows companies to develop new designs that add value to their products and increase productivity. The auto industry is such an example, where (for better or worse) companies realized that customers would pay premium prices for SUVs and minivans. By shifting the mix of their output, auto makers were able to increase their value-added per employee.

Let me clarify here. Clearly, IT has become part of the backbone of business operations, so it does have a role in all of the preceding examples; e.g., in the case of product design, computer-aided designs are important, so IT was certainly a facilitating technology for the innovations. But the key point is that the source of the productivity increase was not primarily major new IT investments, nor, I suspect, were these innovations driven to any great extent by the decline in computer prices.

In general, improvements in workplace organization that affect the way offices are run, the way work flows are organized, and the way workplaces are configured can often have very large payoffs in productivity; and they do seem to have done so, based on the industry case studies. Is there any unifying catalyst for such innovations? The one thing that comes through consistently is that competitive intensity puts a lot of pressure on companies to innovate and improve their productivity. One McKinsey Global Institute (2001) study suggested that this explains the acceleration in productivity that took place in the 1990s. That is a tough case to make, because there is no smoking gun that says competitive intensity suddenly strengthened in the mid-1990s. A better case to be made is that there was a lot of deregulation in the U.S. starting in the 1970s and continuing into the 1980s and 1990s that steadily increased competitive intensity. In addition, there has been a substantial increase in the volume of trade and the extent of globalization in the last twenty years, with a large volume of foreign direct investment coming to the U.S. in the 1990s. In a large country like the U.S., the benefits that are thought to come from globalization often occur as best-practice companies expand their operations throughout the country. It is the progression from regional competition to national competition that increases competitive intensity. In short, there is a pretty good case that competitive intensity in the U.S. economy was greater in the 1990s than in earlier periods, and competitive intensity remains high today.

It may simply be that the forces driving faster productivity growth after the mid-1990s—both IT improvements and increased competitive intensity—were building over a number of years and that business cycle effects or quirks in the data resulted in an apparent pickup in productivity growth starting in the second part of the 1990s. In reality, faster trend labor productivity growth may have been coming down the pipeline before 1995, masked perhaps in the early years of the decade by recession and then by very rapid job growth.

Intangible Capital

There is an emerging body of economics literature looking at investment in intangible capital, which is relevant for understanding productivity trends and for interpreting growth accounting results. Robert Hall, Eric Brynjolfsson, and Robert Gordon have all been leaders in this area. The existence and importance of intangible capital can be used to explain the Solow paradox of why productivity remained slow in the 1980s and early 1990s, even though computer investment was growing rapidly. And it can explain why productivity growth remained rapid after 2000, even though IT spending slumped.

The argument is that it takes time to learn how to apply new technologies effectively and that learning takes resources (people). The buildup in such knowledge is investment in intangible capital, and the resulting knowledge stock is part of an economy's intangible capital. Based on this view, the explanation for the Solow paradox is that, even though companies were investing in IT, measured productivity growth was not increasing because companies also were employing many people to figure out how to use the new technology and how to reorganize their companies to take advantage of the new technology. It was not until after 1995 that the payoff in increased productivity in business operations was large enough to offset the drag on measured productivity from the employees whose output was intangible capital investment. One way of looking at this is that there was (and still is) a measurement error. Spending on tangible capital is counted as investment in GDP, while spending on intangible capital is not.

After 2000, the effects went in the other direction. Under pressure to sustain profits, companies cut back drastically on their investment in IT and cut back similarly on their investment in intangible capital. Measured productivity soared as companies reduced employment, but this came at the price of reduced intangible investment. If this story is correct, then the prospects for future productivity growth may have been compromised. Companies are taking advantage of their past investments and are not building the basis for future growth.

This is an interesting and valuable story, and it can help explain some of the surprising shifts in productivity trends. In some ways it is quite consistent with the industry case study results described earlier. Productivity growth is not coming just from how many computers a company bought this year, or even the stock of computer capital. It also comes from how well companies use information technology. Further, intangible capital does not have to be associated only with IT. Investment in designing new business systems or reorganizing the workplace or developing new products or new markets all can be considered investments in intangible capital. For individual companies, investments in advertising and brand development are important sources of intangible capital. How these contribute to aggregate productivity is less clear.

There is no question that intangible capital is important. That part of the new thinking is solid.[3] The only question about this view is whether the new theory is being presented in quite the right way. One of the lessons from the case studies is that the companies that went out and bought IT and then sat down to figure out how to use it often did not succeed. The innovation should be the starting point, that is, when someone figures out how to improve operations or add value to a company's products. The IT investment that pays off is the investment that follows that innovation. Maybe this nuance would lead to the same place empirically, but it is important to get the story straight. In addition, there is skepticism among the business consultants that I talk to about whether the cutbacks after 2000 really involved the elimination of valuable intangible capital accumulation. The alternative view is that, in the downturn, companies eliminated low-return activities that probably should have been cut earlier (at least from the viewpoint of shareholders). Either way, there has been a one-off temporary surge in productivity, but this alternative view suggests this has not come at the expense of valuable investments in future productivity growth.

Lessons from Other Countries

In looking at what we have learned from the experiences of other countries, I am going to concentrate on Europe,[4] and the story for this region has been changing. There was a straightforward view a couple of years ago that said that productivity growth had not sped up in Europe after 1995; indeed, it may have slowed. At the same time, European businesses apparently had not invested in IT to the same extent as businesses in the U.S. These facts were very supportive of the hypothesis that IT investment had driven faster productivity growth in the U.S. and suggested that Europe needed to increase its IT investment to get on the same bandwagon.

Then new results emerged suggesting that, in fact, Europe had invested rather heavily in IT. Previous studies had underestimated the extent of European IT investment partly because the structure of that investment is different—a lot of in-house software development in Europe was not being reflected in IT spending data. In addition, not all European countries use price deflators that are comparable to those used in the U.S., which again resulted in an understatement of IT investment. Even after allowing for these factors, IT capital stock in Europe still is lower than in the U.S., and certainly IT production is lower. But the gap in IT use is not nearly as large as had been thought.

It is possible that, with further study and revised measurement, the picture we have of Europe will change again. But, based on my assessment of the current state of knowledge, what has happened in Europe is that businesses have greatly increased their use of IT; unlike the U.S., however, the European economies have experienced continued slow productivity growth (with some exceptions).[5] These findings indicate that rapid IT investment growth does not ensure rapid labor productivity growth.

Case studies of industries in France and Germany add to this picture.[6] If you look at the industry in these two economies where productivity grew fastest—mobile telecommunications—the gains are indeed related to IT. The technology became available, it became cheap enough to spread, and the industry grew very rapidly. Indeed mobile telecom productivity grew much faster in France than in the U.S. For that sector, measured labor productivity in France in 2000 was twice as high as it was in the U.S. This is not because the technology in France is particularly better. It is because the regulation of mobile telecom is better in France, which is unusual. In France, they have had about the right amount of competition in mobile telecom. In the U.S. in the 1990s there were too many small companies operating well below minimum efficient scale.

In most of the other case studies, IT did not play a major role in explaining differences in productivity growth among France, Germany, and the U.S. In retailing, productivity is not as high in Germany as it is in the U.S., even though many German retailers use IT in ways that are similar to businesses in the U.S. The difference in performance in retailing really has much more to do with the more limited access to land in Germany and the consequent inability to build big box stores.

Another example is where productivity in the automobile industry grew very fast in France, but not in Germany. That was because in France they were willing to undergo restructuring. In France, they decided to privatize the firms that had been state-owned, and this resulted in changes in management and an acceptance of the need for layoffs. Germany, in contrast, did not go through a similar restructuring; indeed, the German firms face restrictions on lay-

3. Of course the importance of intangible capital in the form of R&D investment has been appreciated for a long time.

4. See Baily and Kirkegaard (2004).

5. Recent work by the OECD (2003) has suggested that differences in measurement methods do not greatly change the standard results for aggregate labor productivity comparisons. These differences may change the industry distribution of growth.

6. See McKinsey (2002b).

offs. IT certainly is important to automobile manufacturers, whether in designing new cars or in using the computer-controlled assembly devices, but differences in the use of IT did not seem to account for differences in productivity growth between the countries in the 1990s.

Electric power is a case where Germany did pretty well in raising its productivity, though not necessarily through great investment in IT. Rather, it did so by privatizing the industry; and even before this occurred, the companies had been restructured in preparation for sale. Germany also changed the structure of regulation in the electricity industry. From the experience of California, we know that regulating electric power is not the easiest thing to do. However, if it is done correctly, there are tremendous gains from privatization and the right kind of regulation. That was the case in Germany, and it seems to have been the case for the United Kingdom also, after some false starts.

In summary, when industries are competitive and not overregulated in Europe, they use IT in ways that are similar to the same industries in the U.S. Differences in IT use came largely as a result of differences in industry structure and regulation, but were not, in any case, the main reason for productivity performance differences. I note that when these case studies were at an early stage, both the McKinsey consultants who did the research and I myself firmly believed that IT would explain much of the performance differences among the countries. As the work progressed, however, it became clear that the case study evidence would not support this view.

Conclusion

Let me bring the focus back to the acceleration of productivity in the U.S. Since the mid-1990s, the acceleration of productivity in the industries that produce IT hardware and the rapid investment in IT capital both contributed to the overall productivity acceleration. The need for businesses to build intangible capital to realize the potential from new technologies is important in understanding the pattern of productivity growth for the U.S. in recent years. Nevertheless, reliance on growth accounting likely has led to an overstatement of the impact of IT.

Productivity is driven by innovation, which may be strongly related to IT use, but often is not. A high level of competitive intensity is important in encouraging leading companies to innovate and in forcing competitors either to adopt the same innovations or to find alternative innovations if they are to survive.

References

Baily, Martin N., and Jacob Kirkegaard. 2004. *A Radical Transformation of the European Economy.* Washington, DC: Institute for International Economics.

Baily, Martin N., and Robert Z. Lawrence. 2001. "Do We Have a New E-Conomy?" *American Economic Review* 91 (May) pp. 308–312.

Foster, Lucia, John Haltiwanger, and C.J. Krizan. 2002. "The Link between Aggregate and Micro Productivity Growth: Evidence from Retail Trade." NBER Working Paper w9120 (August).

Gordon, Robert J. 2000. "Does the 'New Economy' Measure Up to the Great Inventions of the Past?" *Journal of Economic Perspectives* 14(4) (Fall) pp. 49–74.

Jorgensen, Dale W., Man S. Ho, and Kevin J. Stiroh. 2004. "Growth of U.S. Industries and Investments in Information Technology and Higher Education." In *Measuring Capital in a New Economy,* eds. C. Corrado, C. Hulten, and D. Sichel. Chicago: University of Chicago Press.

McKinsey Global Institute. 2001. "U.S. Productivity Growth, 1995–2000." *MGI Reports* (October). www.mckinsey.com/knowledge/mgi/productivity/

McKinsey Global Institute. 2002a. "How IT Enables Productivity Growth." *MGI Reports* (November). www.mckinsey.com/knowledge/mgi/IT/

McKinsey Global Institute. 2002b. "Reaching Higher Productivity Growth in France and Germany." *MGI Reports* (October). www.mckinsey.com/knowledge/mgi/europe/

OECD. 2003. *The Sources of Economic Growth in OECD Countries.* Paris: OECD.

Solow, Robert. 1987. *New York Review of Books* (July 12).

Stiroh, Kevin J. 2002. "Information Technology and the U.S. Productivity Revival: What Do the Industry Data Say?" *American Economic Review* 92(5) (December) pp. 1,559–1,576.

Triplett, Jack E., and Barry P. Bosworth. Forthcoming. "Baumol's Disease Has Been Cured: IT and Multifactor Productivity in U.S. Services Industries." *The New Economy. How New? How Resilient?* ed. Dennis W. Jansen. Chicago: University of Chicago Press.