Volume Title: NBER Macroeconomics Annual 2002, Volume 17

Volume Author/Editor: Mark Gertler and Kenneth Rogoff, editors

Volume Publisher: MIT Press

Volume ISBN: 0-262-07246-7

Volume URL: http://www.nber.org/books/gert03-1

Conference Date: April 5-6, 2002

Publication Date: January 2003

Title: Productivity Growth in the 2000s

Author: J. Bradford DeLong

URL: http://www.nber.org/chapters/c11074

J. Bradford DeLong
UNIVERSITY OF CALIFORNIA AT BERKELEY AND NBER

Productivity Growth in the 2000s

1. *Introduction*

In the early 1970s, U.S. productivity growth fell off a cliff. Measured output per person-hour had averaged a growth rate of 2.8% per year from 1947 to 1973. It averaged a growth rate of only 1.3% per year from 1973 to 1995. In the second half of the 1990s American productivity growth resumed its pre-1973 pace. Between 1995 and the third quarter of 2002, U.S. measured nonfarm-business output per person-hour worked appeared to grow at an annual rate of 2.8% per year.

Nearly all observers agree on the causes of the productivity speedup of 1995–2002. It is the result of the extraordinary wave of technological innovation in computer and communications equipment. Assume that this near-consensus is correct: that the productivity growth speedup in the second half of the 1990s was the result of the technological revolutions in data processing and data communications. What, then, will the future hold? Will the decade of the 2000s see labor productivity growth more like the fast growth seen in the late 1990s? Or more like the slow growth of the 1980s?

In my view, the way to bet is that the next ten years or so will see labor productivity growth as fast as or faster than the U.S. economy has seen since 1995. The answer to the question, "What can we expect from productivity growth in America over the next 10 to 20 years?" is "We can expect very good things."

The case for this point of view follows almost immediately from simple growth accounting and growth theory. The main argument of this paper

I would like to thank the University of California at Berkeley for financial support, and Martin Baily, Susanto Basu, Tim Bresnehan, Nick Crafts, Paul David, Roger Ferguson, John Fernald, Michael Froomkin, Mark Gertler, Robert Gordon, Ned Gramlich, Christopher Gust, Robert Hall, Chad Jones, Boyan Jovanovic, Jaime Marquez, Ken Rogoff, Paul Romer, Suzanne Scotchmer, Andrei Shleifer, Dan Sichel, Larry Summers, Manuel Trajtenberg, Hal Varian, Robert Waldmann, and Gavin Wright for helpful discussions and comments.

begins, after a brief review of the recent history and recent assessments of the cases of changes in productivity growth, by considering the simplest possible growth-theory model that has traction on the major issues. In the near-consensus analysis, increased total factor productivity (TFP) in the information technology (IT) capital-goods-producing sector, coupled with extraordinary real capital deepening as the quantity of real investment in IT capital bought by a dollar of nominal savings grows, has driven the productivity growth acceleration of the later 1990s. The extraordinary pace of invention and innovation in the IT sector has generated real price declines of between 10% and 20% per year in information processing and communications equipment for nearly forty years so far. These extraordinary cost declines have made a unit of real investment in computer or communications equipment absurdly cheap, and hence made the quantity of real investment and thus capital deepening in IT capital absurdly large.[1]

In the 1990s the expanding role and influence of these leading sectors became macroeconomically significant. In a standard growth-accounting framework, the later 1990s saw rapid labor productivity growth because of (1) rapid technological progress in the leading sectors, (2) a healthy share of expenditure on the products of these leading sectors, raising the real IT capital–output ratio, and (3) continued utility of IT capital in production. Continued declines in the prices of IT capital mean that a constant nominal flow of savings channeled to such investments brings more and more real investment. The social return to IT investment would have to suddenly drop to nearly zero, the share of nominal investment spending devoted to IT capital would have to collapse, or technological progress in IT would have to slow drastically, for labor productivity growth in the next decade to reverse itself and return to its late 1970s or 1980s levels. Yet there are no technological reasons for the pace of productivity increase in our economy's leading sectors to decline over the next decade or so. Thus if nominal shares of expenditure on IT capital and of income attributable to existing IT capital remain constant, we can expect the next decade or more to be like the past since 1995. That is the lesson that growth-theory finger exercises have to teach us.

Second, there are four unknown cards in the hole—four reasons to think that the future is likely to be brighter than the simplest models suggest. First, the elasticity of demand for IT goods is likely to remain high. A high elasticity of demand for IT technology goods means that as prices fall, expenditure shares will not remain constant but will rise, boosting

1. It may indeed be the case that a unit of real investment in computer or communications equipment "earned the same rate of return" as any other unit of real investment, as Robert Gordon (2002) puts it. But the extraordinary cheapness of the real unit of capital contributed a major component to the acceleration of labor productivity growth.

growth. Second, the long-time trend shows the share of national income attributable to the returns on the existing IT capital stock rising as well. It would be surprising if this forty-year trend suddenly stopped today. A rising share of national income attributable to existing IT capital would boost growth as well.

Third, Basu, Fernald, and Shapiro (2002) take adjustment costs in investment seriously, and conclude that late-1990s growth undershot trend growth by about 0.5 percentage points per year. Fourth, Paul David (1991) has argued for decades that general-purpose technologies boost labor productivity in two stages: (1) first, capital deepening; (2) second, social learning about how to use new technologies efficiently, a process that drives rapid TFP growth for an extensive period of time but that cannot begin until diffusion is nearly complete. If the pattern he believes holds turns out to hold for IT we can expect to see rapid TFP growth in IT-using industries emerge at some point as an additional growth-promoting factor.

Now it is not likely that all of these hole cards will turn out to be valuable face cards. But it is highly unlikely that none of them will be winners. Thus standard growth-accounting analyses predict a future like the recent past to be a lower bound to reasonable forecasts of future productivity growth.

Is there any reason to be pessimistic? I can think of only one possible hole card on the pessimistic side—a fear of large-scale governmental failure in setting forth the institutional framework to support information-age markets. If governments fail to properly structure the micro marketplace to encourage the growth of high-productivity IT-based industries and practices, then and only then will optimistic macro conclusions be cast into doubt.

2. *The Pattern of Growth in the Later 1990s*

2.1 ASSESSMENTS OF THE RECENT PRODUCTIVITY GROWTH SPEEDUP

In the early 1970s, U.S. productivity growth fell off a cliff. Measured output per person-hour worked in nonfarm business had averaged a growth rate of 2.84% per year from 1947 to 1973. It averaged a growth rate of only 1.34% per year from 1973 to 1995.[2] The productivity slowdown meant

2. The deceleration in the growth rate of total real GDP was somewhat smaller: the social changes that brought more women into the paid labor force in enormous numbers cushioned the effect of this productivity slowdown on the growth rate of measured total real GDP, if not its effect on Americans' material welfare.

that, according to official statistics, Americans in 1995 were only 70% as productive as their predecessors back in the early 1970s would have expected them to be. The productivity slowdown gave rise to an age of diminished expectations that had powerful although still debated effects on American politics and society.[3]

In the second half of the 1990s American productivity stood up, picked up its mat, and walked. It resumed its pre-1973 pace. Between 1995 and the third quarter of 2002, U.S. measured nonfarm-business output per person-hour worked appeared to grow at an annual rate of 2.79% per year.[4]

Noneconomists tended to attribute a large chunk of fast late-1990s growth to "cyclical" factors.[5] Economists, however, have had a much harder time attributing more than a few tenths of a percentage point per year of late-1990s growth to the business cycle.[6] The standard indicators of high cyclically-driven productivity were absent. Moreover, as Susanto Basu, John Fernald, and Matthew Shapiro have argued, there are stronger reasons for thinking that the adjustment costs associated with moving to a more IT-capital-intensive growth path led actual growth to understate trend growth than for thinking that cyclical factors led actual growth to overstate trend growth in the second half of the 1990s.[7] And the extremely rapid runup of stock prices indicated that at least the marginal investor

3. See Krugman (1994) for one interpretation of how the productivity slowdown made a significant difference.
4. Figuring out what the growth rate of real output has been since 1994 poses unusual challenges. The most important of these is the discrepancy between national product and national income. In 1994 the statistical discrepancy between the two—the amount you had to add to national product in order to get to national income after making all of the conceptual and definitional adjustments—was +$59 billion. By 2000 this statistical discrepancy was −$130 billion. National income grew by an extra $190 billion relative to national product between 1994 and 2000, not because of conceptual definitions but because of errors and omissions (or, rather, inconsistent and changing patterns of errors and omissions). By now this shift in the statistical discrepancy adds up to an amount equal to 3.5% of national product. Take national product as your guide to the growth of the American economy, and you conclude that measured real labor productivity growth from 1995 to 2002 was 2.63% per year. Take national income as your guide and you conclude that it was 2.95% per year. Split the difference—as Martin Baily (2002) recommends—and you conclude that it was 2.77% per year. I am agnostic as to which of the two measures of the economy's size is going awry, and so I follow Baily and split the difference. Note that this divergence between product- and income-side measures of economic output is very recent. Between 1959 and 1973, product grew faster than income by an average of only 0.04% per year. Between 1973 and 1995, income grew faster than product by an average of 0.02% per year. The sustained growth discrepancy between 1995 and 2002 of 0.35% per year is extremely unusual.
5. See, for example, Kosterlitz (2002).
6. See Gordon (2002) and Gordon (2000a, 2000b).
7. See Basu, Fernald, and Shapiro (2001).

in equities anticipated that the acceleration of economic growth that started in the mid-1990s would last for decades or longer.[8]

The causes of the productivity slowdown of 1973–1995 or so remain disappointingly mysterious. Baily (2002) calls the growth-accounting literature on the slowdown "large but inconclusive." No single factor provides a convincing and coherent explanation, and the residual position that a large number of growth-retarding factors suddenly happened to hit at once is but the least unlikely of the residual explanations.[9]

By contrast, nearly all agree on the causes of the productivity speedup of 1995–2001: it is the result of the extraordinary wave of technological innovation in computer and communications equipment. Even though the failure of economists to reach consensus in their explanations of the productivity slowdown has to leave one wary of the reliability of the consensus about the causes of the productivity speedup, the depth and range of this near-consensus is remarkable.

Robert Gordon (2002) writes that cyclical factors account for "0.40" percentage points of the growth acceleration, and that the rest is fully accounted for by IT—an "0.30 [percentage] point acceleration [from] MFP growth in computer and computer-related semiconductor manufacturing" and a "capital-deepening effect of faster growth in computer capital . . . [that] in the aggregate economy accounts [for] 0.60 percentage points of the acceleration."

Kevin Stiroh (2001) writes that "all of the direct contribution to the post-1995 productivity acceleration can be traced to the industries that either produce [IT capital goods] or use [them] most intensively, with no net contribution from other industries . . . relatively isolated from the [IT] revolution."

Oliner and Sichel (2000) write that "the rapid capital deepening related to information technology capital accounted for nearly half of this increase" in labor productivity growth, with a powerful "additional growth contribution . . . com[ing] through efficiency improvement in the *production* of computing equipment." Jorgenson, Ho, and Stiroh (2001) reach the same conclusions about the importance of IT capital deepening and increased efficiency in the production of computing and communications

8. See Greenwood and Jovanovic (1999).

9. See Fischer (1988), Griliches (1988), Jorgenson (1988), and Gordon (2000b, 2002). Jorgenson (1988) convincingly demonstrates that the oil price shocks can plausibly account for slow growth in potential output in the 1970s, but why does potential output growth remain slow after 1986 after real oil prices have fallen again? Griliches (1988) finds that an explanation in terms of a slowdown in innovation is unattractive, but Gordon (2000b, 2002) finds such an explanation attractive.

equipment as major drivers of the productivity growth acceleration, and they go on to forecast that labor productivity growth will be as high in the next decade as it has been in the past half decade.[10]

The only major empirical study taking a stand against this explanation is that of the McKinsey Global Institute (2001), which presents a regression of the growth in value added per worker and the increase in computer capital by industry. When industry observations are counted equally, it finds next to no correlation between computer capital and labor productivity. When industries are weighted by employment, it finds a statistically significant and substantively important connection. It is unclear why the McKinsey Global Institute prefers its unweighted regressions to its weighted ones.

In its case studies the McKinsey Global Institute attributes rapid growth in productivity in the retail distribution sector to managerial innovations on the part of Wal-Mart, coupled with competitive pressure exerted by Wal-Mart on the rest of the sector. However, Wal-Mart's founders and executives have long attributed much of their competitive success to the skillful and intensive use of IT. For example, Sam Walton (1992) wrote in his autobiography (quoted in Cohen, DeLong, and Zysman, 2000) about how "information sharing [was] a new source of power . . . [W]e believed in showing a store manager every single number relating to his store, and eventually we began sharing those numbers with the department heads in our stores. . . . That's why we've spent hundreds of millions of dollars on computers and satellites—to spread all the little details around the company as fast as possible. But they were worth the cost. It's only because of information technology that our store managers have a really clear sense of what they're doing most of the time. . . ."

One of the few prominent economists who appear to expect slow productivity growth over the medium turn is Joseph Stiglitz. As Stiglitz is quoted by Kosterlitz (2002): "The fact that things have stabilized does not mean they've recovered. When people say things are not so bleak, they mean that the economy is not in free fall, not in a negative spiral." Kosterlitz goes on to write, "The recovery, [Stiglitz] says, might not be as snappy as the conventional models used by the forecasters suggest. 'Most downturns have been inventory recessions. They tend to be short-lived; as companies deplete their inventories, things improve. This is different. It's not just an inventory downturn, but also a case of overcapacity in areas where there was lots of investing—IT, telecom. . . . These represent a significant

10. However, Jorgenson, Ho, and Stiroh expect total real GDP growth to slow because of slower growth in hours worked—they forecast 1.1% per year growth in hours over the next decade, compared to 2.3% per year from 1995 to 2000.

share of investment in the late 1990s,' said Stiglitz. Things won't improve until industry gets rid of excess equipment and employees, he says. 'The real restructuring takes time.'"

Such arguments that recessions are the result of "overinvestment" which must inevitably lead to a period of slow growth during which the overhang of excess capital is "liquidated" have often been made in economics (see DeLong, 1991). But it appears, to me at least, hard to sustain a claim that we are in such a situation today. With the prices of IT goods falling as rapidly as they are, surely real capital-output ratios in IT sectors are below their long-run values. In such a case, it makes no sense to claim that there is a capital overhang to be "liquidated": such a claim requires that the ratio of the real stock of IT capital to output be above its long-run level. Investment in the near future in many IT sectors may be low, but low investment seems much more likely to be attributable to low demand or to failures of appropriability of the products of investment than to too much capital.

2.2 INFORMATION TECHNOLOGY AND POST-W.W. II ECONOMIC GROWTH

Compare our use of IT today with our predecessors' use of IT half a century ago.[11] The decade of the 1950s saw electronic computers largely replace mechanical and electromechanical calculators and sorters as the world's automated calculating devices. By the end of the 1950s there were roughly 2000 installed computers in the world: machines like Remington Rand UNIVACs, IBM 702s, or DEC PDP-1s. The processing power of these machines averaged perhaps 10,000 machine instructions per second.

Today, talking rough orders of magnitude only, there are perhaps 300 million active computers in the world with processing power averaging several hundred million instructions per second. Two thousand computers times 10,000 instructions per second is 20 million. Three hundred million computers times, say, 300 million instructions per second is 90 quadrillion—a 4-billion-fold increase in the world's raw automated computational power in forty years, an average annual rate of growth of 56%.

Such a sustained rate of productivity improvement at such a pace is unprecedented in our history. Moreover, there is every reason to believe that this pace of productivity growth in the leading sectors will continue for decades. More than a generation ago Intel Corporation's co-founder Gordon Moore noticed what has come to be called Moore's law—that improvements in semiconductor fabrication allow manufacturers to dou-

11. For an extended version of this part of the argument, see Cohen, DeLong, and Zysman (2000).

ble the density of transistors on a chip every eighteen months. The scale of investment needed to make Moore's law hold (Figure 1) has grown exponentially along with the density of transistors and circuits, but the law has continued to hold, and engineers see no immediate barriers that will bring the process of improvement to a halt anytime soon.

2.2.1 Investment Spending As the computer revolution proceeded, nominal spending on IT capital rose (Figure 2) from about 1% of GDP in 1960 to about 2% of GDP by 1980 to about 3% of GDP by 1990 to between 5% and 6% of GDP by 2000. All throughout this time, Moore's law meant that the real price of IT capital was falling as well. As the nominal spending share of GDP spent on IT capital grew at a rate of 5% per year, the measured price of information-processing equipment plus software fell steadily at a pace between 5% and 10% per year.

At chain-weighted real values constructed using 1996 as a base year, real investment in IT equipment and software was equal to 1.7% of real GDP in 1987 (although it is important to remember that this does not mean that real investment in IT equipment plus software was a 1.7% share

Figure 1 REAL INVESTMENT IN INFORMATION TECHNOLOGY EQUIPMENT AND SOFTWARE, AND REAL GDP

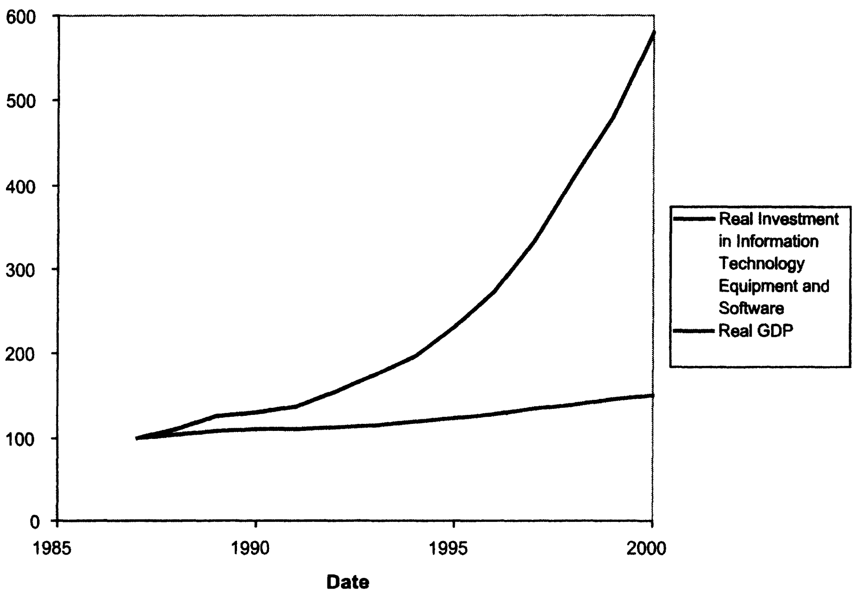

Source: National Income and Product Accounts.

Figure 2 NOMINAL SPENDING ON INFORMATION TECHNOLOGY EQUIPMENT PLUS SOFTWARE AS A SHARE OF NOMINAL GDP

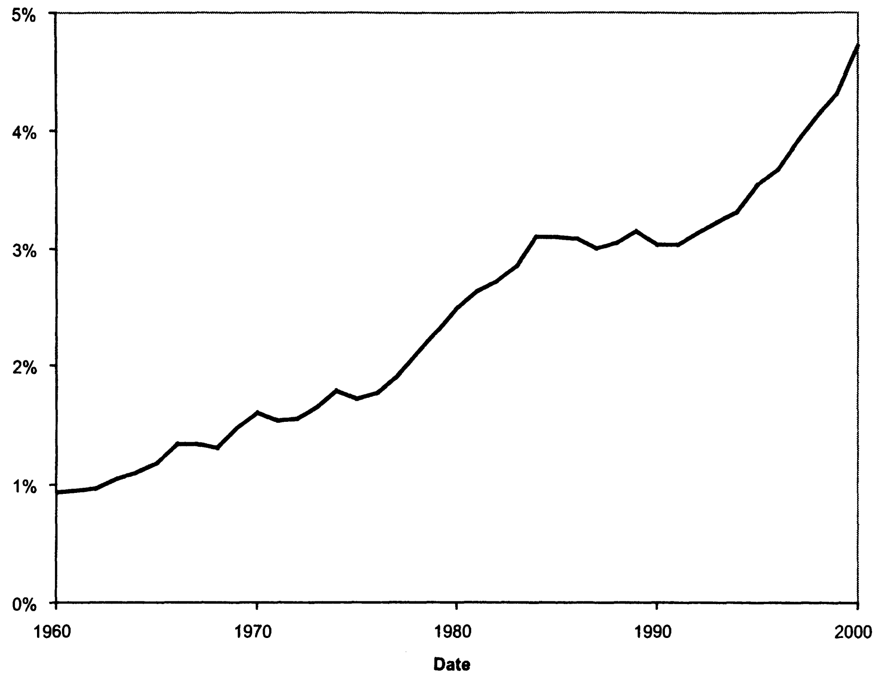

of GDP; in the world of chain-weighted statistics, real components do not sum to their aggregate). By 2000 it was equal to 6.8% of real GDP.[12] The steep rise in real investment in information-processing equipment (and software) drove a steep rise in total real investment in equipment: by and large, the boom in real investment in information-processing equipment driven by rapid technological progress and the associated price declines was an addition to, not a shift in, the composition of overall real equipment investment.

2.2.2 Macro Consequences A naive back-of-the-envelope calculation would suggest that this sharp rise in equipment investment was of sufficient magnitude to drive substantial productivity acceleration: at a total social rate of return to investment of 15% per year, a 6-percentage-point rise in the investment share would be predicted to boost the rate of growth of real gross product by at least about 1 percentage point per year. And

12. For an excellent overview of what forms of addition and comparison are or are not legitimate using real chain-weighted values, see Whelan (2000a).

that is the same order of magnitude as the 1.0- to 1.6-percentage-point acceleration in annual labor productivity growth rates seen in the second half of the 1990s.

The acceleration in the growth rate of labor productivity (Figure 3) and of real GDP in the second half of the 1990s effectively wiped out all the effects of the post-1973 productivity slowdown. The U.S. economy in the second half of the 1990s was, according to official statistics and measurements, performing as well in terms of economic growth as it had routinely performed in the first post-W.W. II generation. It is a marker of how much expectations had been changed by the 1973–1995 period of slow growth that 1995–2001 growth was viewed as extraordinary and remarkable.

Nevertheless, the acceleration of growth in the second half of the 1990s was large enough to leave a large mark on the economy even in the relatively short time it has been in effect. Real output per person-hour worked in the nonfarm business sector today is 10% higher than one would have predicted back in 1995 by extrapolating the 1973–1995 trend. That such a large increase in the average level of productivity can be accumulated

Figure 3 LOG LABOR PRODUCTIVITY IN NONFARM BUSINESS

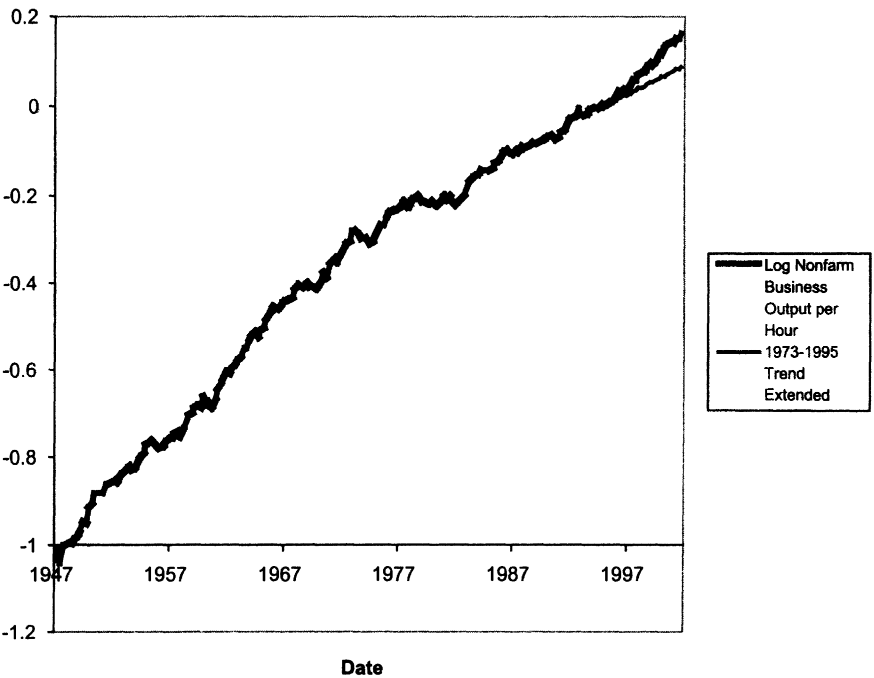

Source: National Income and Product Accounts.

Figure 4 UNEMPLOYMENT AND THE CHANGE IN INFLATION

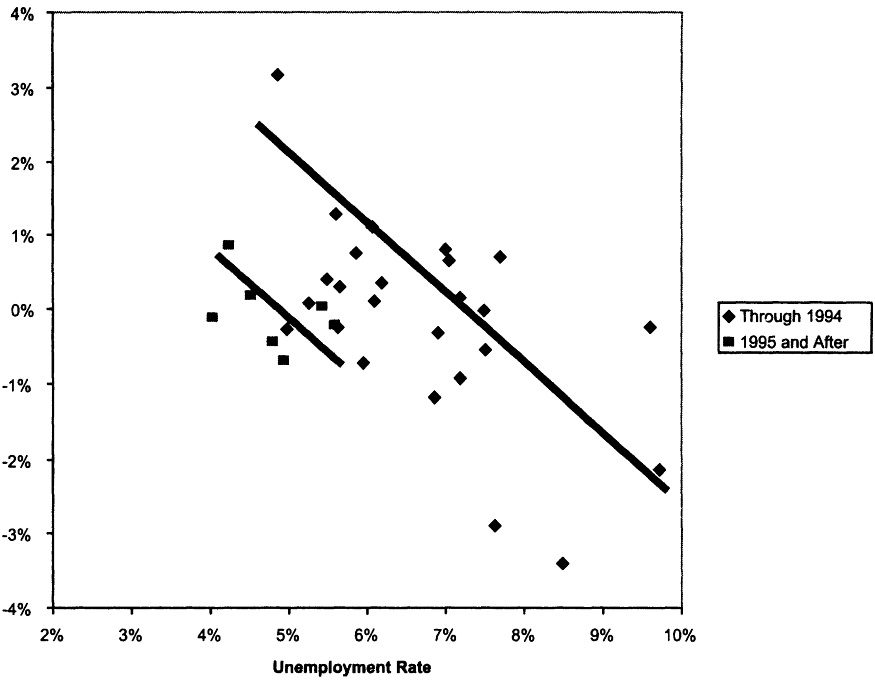

Source: Bureau of Labor Statistics, Bureau of Economic Analysis.

over a mere seven years just by getting back to what seemed "normal" before 1973 is an index of the size and importance of the 1973–1995 productivity slowdown.

2.2.3 Cyclical Factors Alongside the burst of growth in output per person-hour worked came significantly better labor-market performance (Figure 4). The unemployment rate consistent with stable inflation, which had been somewhere between 6% and 7% of the labor force from the early 1980s into the early 1990s, suddenly fell to 5% or even lower in the late 1990s. All estimates of nonaccelerating-inflation rates of unemployment (NAIRUs) are hazardous and uncertain,[13] but long before 2001 the chance that the inflation–unemployment process was a series of random draws from the same urn after as before 1995 was negligible.

This large downward shift in the NAIRU posed significant problems for anyone wishing to estimate the growth of the economy's productive

13. See Staiger, Stock, and Watson (1997).

potential over the 1990s. Was this fall in the NAIRU a permanent shift that raised the economy's level of potential output? Was it a transitory result of good news on the supply-shock front—falling rates of increase in medical costs, falling oil prices, falling other import prices, and so forth—that would soon be reversed? If the fall in the NAIRU was permanent, then presumably it produced a once-and-for-all jump in the level of potential output, not an acceleration of the growth rate of potential output. But how large a once-and-for-all jump? Okun's law would suggest that a 2-percentage-point decline in the unemployment rate would be associated with a 5% increase in output.

Production-funtion-based analyses, however, would suggest that a 2-percentage-point decline in the unemployment rate would be associated with a roughly 1.5% increase in output. One must take account of the effect of falling unemployment on the labor force and the differential impact of the change in unemployment on the skilled and the educated.

However, none of the other available cyclical indicators suggest that the late-1990s economy was an unusually high-pressure economy. The average workweek (Figure 5) was no higher in 2000, when the unemployment rate approached 4% than it had been in 1993, when the unemployment rate fluctuated between 6% and 7%.

Capacity utilization (Figure 6) was lower during the late 1990s than it had been during the late 1980s, when unemployment had been 1.5 percentage points higher.[14] Low and not rising inflation, a relatively short workweek, and relatively low capacity utilization—these all suggested that the fall in the unemployment rate in the late 1990s was not associated with the kind of high-pressure economy assumed by Okun's law.

3. *A Simple Model*

Given that the acceleration in productivity growth in the second half of the 1990s was primarily driven by the revolution in IT, what conclusions can be drawn about the likely pace of productivity growth in the future? The first step in answering this question is to write down a simple model that has at least some traction on the major issues. The second step is then to use that model to analyze future productivity growth. And the third step is to step back from the model, and to consider the importance of the factors that the model leaves out.

The simple model will be one in which there is (1) an ongoing technological revolution in the production of data-processing and data-commu-

14. One reason, however, for the low measured capacity utilization in the late 1990s was the belief that high levels of investment were expanding capacity at a furious rate.

Figure 5 AVERAGE BUSINESS WORKWEEK

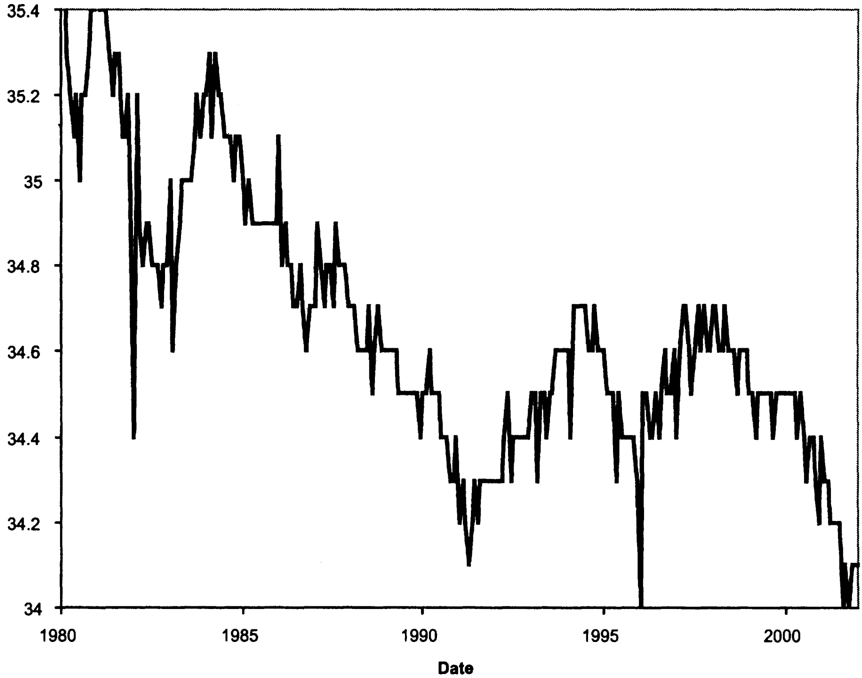

Source: Bureau of Labor Statistics.

nications capital, which is (2) an important input into production. The analysis of the model will largely turn on whether there is reason to anticipate that the pace of technological progress in IT is going to slow, that the share of GDP spent on IT investment is going to decline, or that the share of national product attributable to returns earned on the existing IT capital stock is going to fall. Yet if none of the three of these happens, then this growth-theory finger exercise predicts that we can expect the medium-run future to be as bright for measured productivity growth as the recent past has been.

The third stage consists of analyzing the four hole cards—four reasons to think that the future is likely to be brighter than the simplest possible model suggests. First, the elasticity of demand for IT is likely to remain high—which means that as prices fall, expenditure shares will not remain constant but will rise, boosting growth. Second, the long-time trend shows the income share attributable to IT capital rising: it would be surprising if this forty-year trend suddenly stopped today, and rising income shares of IT boost growth as well. Third, Basu, Fernald, and Shapiro take adjust-

Figure 6 CAPACITY UTILIZATION (%)

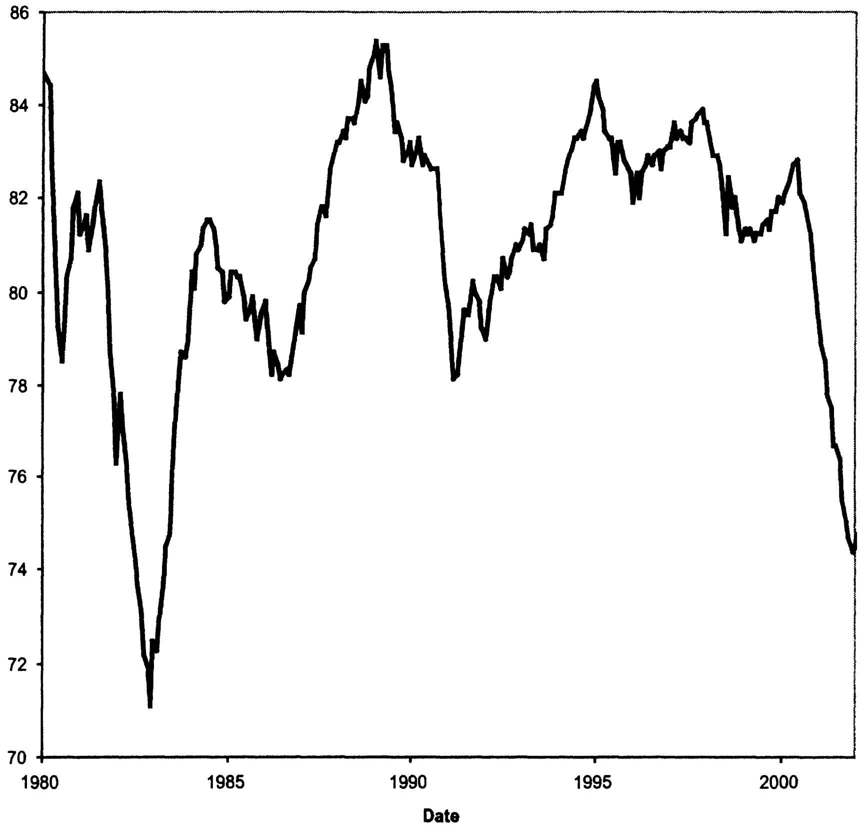

Source: Federal Reserve.

ment costs in investment seriously, and conclude that late-1990s growth undershot trend growth by about 0.5 percentage points per year. Fourth, Paul David has argued for decades that general-purpose technologies boost labor productivity in two stages: (1) first, capital deepening; (2) second, social learning about how to use new technologies efficiently, a process that drives rapid TFP growth for an extensive period of time but that cannot begin until diffusion is nearly complete.

Surely not all of these will turn out to be valuable face cards. But I think it highly unlikely that none of them will win any hands. And I think they are likely to outweigh the only possible hole card on the pessimistic side that I can think of—large-scale governmental failure in setting forth the institutional framework to support information-age markets.

3.1 BASIC THEORY

Suppose that the economy produces two types of output—regular goods, which we will denote by a superscript r, and IT capital, which we will denote by an i. At each moment in time there is a set cost price p_t^i at which output in the form of regular goods can be transformed into IT capital and vice versa.

Let the economy produce two types of output—regular goods, which we will denote by an r, and IT capital, which we will denote by an i. And let at each moment in time the relative price of IT goods be a constant p_t^i, at which output in the form of regular goods can be traded off into IT capital and vice versa. Thus

$$Y_t = Y_t^r + p_t^i Y_t^i. \tag{1}$$

With regular goods serving as numeraire, the total output Y_t is equal to the output Y_t^r of regular goods plus $p_t^i Y_t^i$, the output of IT capital multiplied by its current cost price.

The total output Y_t is itself determined by a standard production function:

$$Y_t = F(A_t, K_t^r, K_t^i, L_t), \tag{2}$$

where A_t is the exogenous level of TFP, K_t^r is the stock of *normal* (non-IT) capital, K_t^i is the stock of IT capital, and L_t is the labor force.

Suppose further that because of ongoing technological revolutions the cost price of IT is declining at a constant proportional rate π. Now note that Y_t is not a measure of real output—real output will grow faster than Y_t as p_t^i falls, because the ability to make IT goods more cheaply is a source of productivity growth as well. I will use Y_t^* to stand for a chain-weighted measure of real output so that we capture this additional source of growth.

Then in this framework the proportional rate of growth of real chain-weighted real output Y^* will be

$$\frac{d \ln Y_t^*}{dt} = \frac{\partial F}{\partial A}\frac{A}{Y}\frac{d \ln A_t}{dt} + \frac{\partial F}{\partial K^r}\frac{K_t^r}{Y}\frac{d \ln K_t^r}{dt} + \frac{\partial F}{\partial K^i}\frac{K_t^i}{Y}\frac{d \ln K_t^i}{dt} + \frac{\partial F}{\partial L}\frac{L}{Y}\frac{d \ln L_t}{dt} + X_t^i \pi. \tag{3}$$

The rate of growth of real output will be equal to contributions from labor, normal capital, IT capital, and TFP in the production of regular output, plus an extra term equal to the share of total expenditure on IT capital X_t^i times the rate π at which the cost price of IT goods is declining.

Under the assumptions of constant returns to scale and competition, $(\partial F/\partial K)(K/Y)$ and like terms are simply the shares of national income appropriated by each of the three factors of production. So let us use s_i, s_r, and s_L as a shorthand for those terms, and also normalize the TFP term A, and thus rewrite (3) as

$$\frac{d \ln Y_t^*}{dt} = \frac{d \ln A_t}{dt} + s_r \frac{d \ln K_t^r}{dt} + s_i \frac{d \ln K_t^i}{dt} + s_L \frac{d \ln L_t}{dt} + X_t^i \pi. \tag{4}$$

If we assume a constant proportional growth rate n for the labor force, a constant growth rate a for TFP in the production of normal output Y, and constant shares of nominal expenditure X^r and X^i on normal and IT gross investment, then (4) becomes

$$\frac{d \ln Y_t^*}{dt} = a + s_L n + s_r\left(\frac{X^r Y}{K^r} - \delta^r\right) + s_i\left(\frac{X^i Y}{p_t^i K^i} - \delta^i\right) + X^i \pi. \tag{5}$$

And if we are willing to impose constant returns to scale in the three factors of labor, normal capital, and IT capital, then we can rewrite (5) with the rate of growth of labor productivity on the left-hand side as

$$\frac{d \ln(Y_t^*/L_t)}{dt} = a + s_r\left(\frac{X^r Y}{K^r} - (\delta^r + n)\right) + s_i\left(\frac{X^i Y}{p_t^i K^i} - (\delta^i + n)\right) + X^i \pi. \tag{6}$$

3.2 IMPLICATIONS OF BASIC THEORY

The first two terms on the right-hand side are very standard: the TFP growth a, and the contribution from the deepening of the ratio of normal capital per worker:

$$s_r\left(\frac{X^r Y}{K^r} - (\delta^r + n)\right), \tag{7}$$

equal to the normal *capital share* s_r times the net proportional rate of growth of the normal capital stock—its expenditure share X^r divided by the capital–output ratio K^r/Y, minus the labor-force growth rate n plus the depreciation rate δ^r.

But there are the two extra terms. The second term,

$$X^i \pi, \tag{8}$$

is what Oliner and Sichel refer to as the "additional growth contribution . . . com[ing] through efficiency improvement in the *production* of computing equipment." Even if the level of potential normal output were to remain constant, the fact that the economy is able to make IT capital more and more cheaply in terms of normal goods is a genuine improvement in productivity.

The first term,

$$s_i\left(\frac{(X^i/P_t^i)Y}{K^i} - (\delta^i + n)\right), \tag{9}$$

is the contribution to the production of normal output from the net increase in IT capital stock per worker. However the numerator is not the nominal share of GDP expended on IT capital X^i, but the real share X^i/p_t^i. And—because the cost price of IT capital is falling at the rate π—a constant nominal expenditure share means that the real expenditure share relevant for the contribution of IT capital to output growth is growing at a proportional rate π. It is no surprise at all that as long as the nominal expenditure share on IT capital remains constant and the technological revolution is ongoing, the economy exhibits a steadily rising real gross investment expenditure share X^i/p_t^i, and a steadily rising ratio of real IT capital to normal output.[15]

This is in fact what happened in the original industrial revolution: as the dynamic modern sector grew to encompass the bulk of the economy, overall productivity growth accelerated.[16] The heroic age of double-digit annual productivity increase within the steam-power and textile-spinning sectors of the economy ended before the nineteenth century was a quarter over. Yet the major contribution of steam power and textile machinery to British aggregate economic growth took place in the middle half of the nineteenth century. Thus historians of the British industrial revolution like Landes (1969) focus on the late eighteenth century, while macroeconomists and sociologists focus on the mid-nineteenth century: the lag in time between the major innovations and fastest proportional growth of the leading sector on the one hand, and its major influence on aggregates on the other, is likely to be substantial.

If we follow Whelan (2001) and define as auxiliary variables the nomi-

15. There are some subtleties about what is the right way to measure output and how to define a "steady state" in models like this. Exactly what is the most useful way is insightfully explored by Whelan (2001).

16. See Crafts (1985).

nal regular capital–output ratio κ_t^r and the nominal current-dollar-value IT capital–output ratio κ_t^r by

$$\kappa_t^r = \frac{K_t^r}{Y_t}, \tag{10}$$

$$\kappa_t^i = \frac{p_t^i K_t^i}{Y_t}, \tag{11}$$

then we can construct a pseudo-steady-state path for this economy. In the equation for the proportional rate of change of regular output Y,

$$\frac{d \ln(Y_t/L_t)}{dt} = a + s_r\left(\frac{X^r Y_t}{K_t^r} - (\delta^r + n)\right) + s_i\left(\frac{X^i Y_t}{p_t^i K_t^i} - (\delta^i + n)\right), \tag{12}$$

we can substitute in these auxiliary nominal capital–output ratios:

$$\frac{d \ln(Y_t/L_t)}{dt} = a + s_r\left(\frac{X^r}{\kappa_t^r} - (\delta^r + n)\right) + s_i\left(\frac{X^i}{\kappa_t^i} - (\delta^i + n)\right), \tag{13}$$

and then derive rates of change of these ratios:

$$\frac{d\kappa_t^r}{dt} = (1 - s_r)X^r - [a + (1 - s_r)(\delta^r + n) - s_i(\delta^i + n)]\kappa_t^r - s_i X^i\left(\frac{\kappa_t^r}{\kappa_t^i}\right), \tag{14}$$

$$\frac{d\kappa_t^i}{dt} = (1 - s_i)X^i - [a + \pi + (1 - s_i)(\delta^i + n) - s_r(\delta^r + n)]\kappa_t^i - s_r X^r\left(\frac{\kappa_t^i}{\kappa_t^r}\right). \tag{15}$$

We also substitute the nominal capital–output ratios into the production function:

$$\frac{Y_t}{L_t} = A_t\left(\frac{K_t^r}{L_t}\right)^{s_r}\left(\frac{K_t^i}{L_t}\right)^{s_i}, \tag{16}$$

to obtain

$$\frac{Y_t}{L_t} = A_t^{1/(1-s_r-s_i)}(\kappa_t^r)^{s_r/(1-s_r-s_i)}(\kappa_t^i)^{s_i/(1-s_r-s_i)}(p_t^i)^{-s_i/(1-s_r-s_i)}. \tag{17}$$

The dynamics of output per worker in the economy can then be analyzed in terms of the (constant) proportional increase in the TFP A, the (constant) proportional decrease in the real cost price of IT goods, and the dynamic evolution of the nominal capital–output ratios:

$$\frac{d\ln(Y_t/L_t)}{dt} = \frac{a}{1-s_r-s_i} + \frac{s_i\pi}{1-s_r-s_i} + \frac{s_r}{1-s_r-s_i}\frac{d\ln\kappa_t^r}{dt} + \frac{s_i}{1-s_r-s_i}\frac{d\ln\kappa_t^i}{dt}. \tag{18}$$

From (14), (15), and (18), we can calculate the behavior of the economy in its long-run pseudo-steady state. We can see that the proportional growth rate of Y/L will be

$$\frac{d\ln(Y_t/L_t)}{dt} = \frac{a}{1-s_r-s_i} + \frac{\pi s_i}{1-s_r-s_i} \tag{19}$$

and the long-run growth rate of real output per worker will be

$$\frac{d\ln(Y_t^*/L_t)}{dt} = \frac{a}{1-s_r-s_i} + \frac{\pi s_i}{1-s_r-s_i} + X^i\pi, \tag{20}$$

which is the sum of three terms: a term capturing the effect of background TFP growth a on the economy, a term capturing the effect of ongoing capital deepening made possible by falling IT capital prices, and a term capturing the direct effect of improvements in efficiency in the production of IT goods.

How to calibrate this simple theoretical result (20) to the American economy? Over the years since 1995, X^i—the share of expenditure on IT and related forms of capital—has averaged some 6% of GDP. According to Oliner and Sichel (2000), s_i—the share of income attributable to IT capital—has averaged 7% of GDP. Assuming an income share for other kinds of capital of 33%, then at a 10%-per-year decline in the real prices of IT goods the last two terms—the acceleration terms—of equation (20) amount to a productivity-growth boost of 1.8% per year. This is a larger boost to growth than was in fact seen in the acceleration in the later 1990s. Thus this growth accounting exercise certainly does not suggest that productivity growth in the next decade will be lower than in the recent past, as long as Moore's law continues to hold and the prices of IT goods continue to decline rapidly.

However, there is no special reason to think that such a steady-state analysis is the best we can do. The steady state assumes constant nominal investment shares in IT capital, a constant rate of real price decrease in this technologically explosive leading sector, and a constant share parameter s_i. Yet all the evidence we have suggests that all three of these variables move, and move radically, in a decade or less. The American economy began the 1980s very far away from its pseudo-steady state: back then the GDP share of nominal spending on IT investment was only 40% of its current value, and likewise for the share of national income attributable to the IT capital stock.

Thus the potential importance of the first hole card: economies that are approaching their steady-state growth paths from below grow faster than in steady state. With two kinds of capital goods depreciating at different

Figure 7 SAMPLE SIMULATIONS: SPEED OF CONVERGENCE TO THE PSEUDO-STEADY-STATE GROWTH PATH

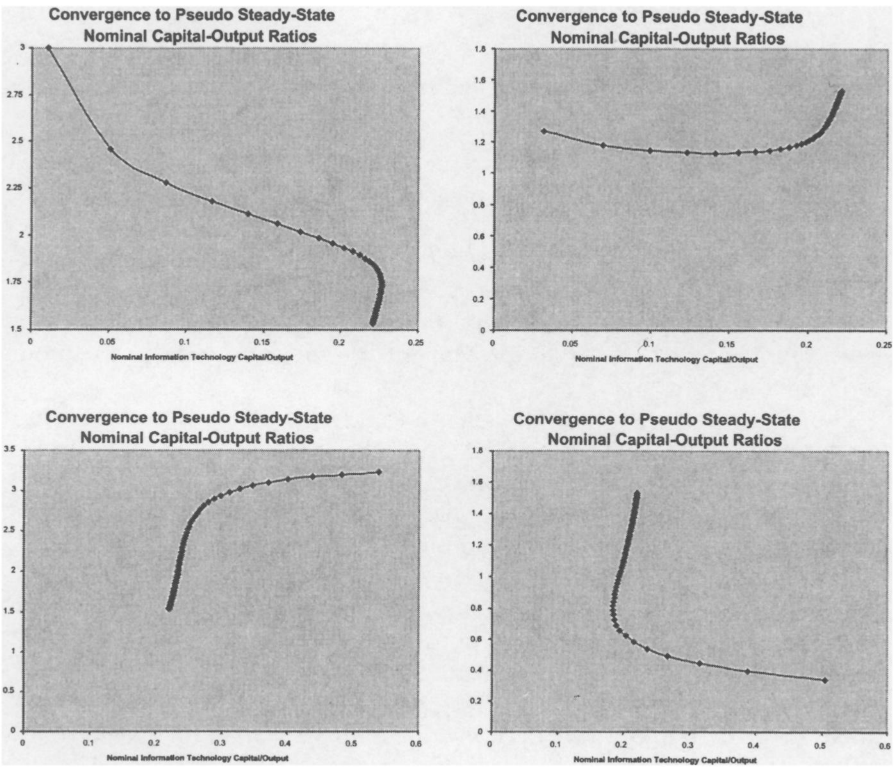

Parameters: $a = 0.01$, $n = 0.01$, $X^i = 0.06$, $X^r = 0.14$, $s_i = 0.1$, $s_r = 0.5$, $\delta^i = 0.1$, $\delta^r = 0.3$, $\pi = 0.1$.

rates, the simple model has no theoretically tractable dynamics. However, simulations (Figure 7) suggest that it is overwhelmingly likely that in such a case the nominal capital–output ratio for the most rapidly depreciating kind of capital converges fairly rapidly—within two decades—to a value near its steady-state value.

This suggests that any additional contribution to growth from convergence dynamics will be confined to the next decade or so.

4. *A Demand-Side Model*

4.1 A MODEL OF CHANGING DEMAND SHARES

However, this leaves unexamined the question of why the income and expenditure shares of IT capital have been rising so rapidly over the past several decades. In order to even grapple with these questions, we have to look on the demand-for-IT side. What determines whether the demand for IT capital and thus for investment grows more rapidly or less rapidly than output as a whole?

An alternative approach is to simplify the production side of the model radically, and instead focus on the implications of changing prices of IT goods for demand. If the TFP in the rest of the economy is growing at a rate π_R, and if the TFP in the leading industries and sectors is growing at a faster rate π_L, then the TFP growth in the economy as a whole will be equal to

$$\pi = \sigma\pi_L + (1 - \sigma)\pi_R, \tag{21}$$

where σ is the share of total expenditure on the goods produced by the economy's fast-growing technologically dynamic leading sectors.

As the process of innovation and technological revolution in the leading sectors proceeds, we would not expect the leading sector share σ of total expenditure to remain constant. If the goods produced by the leading sector have a high (or low) price elasticity of demand, the falls over time in their relative prices will boost (or reduce) the share of total expenditure σ: only if the price elasticity of demand, ε_P, is one will the fall in the relative price of leading-sector products produced by the technological revolutions leave the leading-sector share unchanged.[17]

17. The demand share will also depend on the income elasticity of demand. If the goods produced by the leading sectors are superior (or inferior) goods, the share σ will rise (or fall) as economic growth continues: only if the income elasticity of demand, ε_I, for its products is one will changes in the overall level of prosperity leave the leading sector share unchanged. But I will not model this effect here.

Moreover, the leading-sector share of total expenditure σ matters only as long as the leading sector remains technologically dynamic. Once the heroic phase of invention and innovation comes to an end and the rate of TFP growth returns to the economy's normal background level π_R, the rate of productivity growth in the economy as a whole will return to that same level π_R and the leading-sector share of expenditure σ will no longer be relevant.

Thus four pieces of information are necessary to assess the aggregate economic impact of an explosion of invention and innovation in a leading sector:

The initial share of expenditure on the leading sector's products, σ_0.
The magnitude of the relative pace of cost reduction, $\pi_L - \pi_R$, during the leading sector's heroic age of invention and innovation.
The duration of the leading sector's heroic age of invention and innovation.
The price elasticity of demand, ε_P, for the leading sector's products.

To gain a sense of the importance of these factors, let's consider a few simulations with sample parameter values. For simplicity's sake, set the initial share of expenditure on the leading sector's products, σ_0, equal to 0.02; set the income elasticity of demand for the leading sector's products, ε_I, equal to 1.0; set the heroic age of invention and innovation to a period 40 years long; and set the background level of TFP growth, π_R, to 0.01 per year. Consider three values for the price elasticity of demand, ε_P: 0.5, 2.0, and 4.0. And consider two values for the wedge in the annual rate of technological progress between the leading sector and the rest: 0.03 and 0.05.

With a price elasticity of demand of 0.5, the expenditure share of the leading sectors declines from its original value of 2% as technology advances and the prices of leading-sector goods fall. With a productivity wedge of 5% per year, the initial rate of growth of economy-wide productivity growth is 1.1% per year—1% from the background growth of the rest of the economy, and an extra 0.1% from the faster productivity growth in the one-fiftieth of the economy that is the leading sector. By the twelfth year the expenditure share on leading-sector products has fallen below 1.5%. By the twenty-eighth year it has fallen below 1.0%. By the fortieth year it has fallen to 0.7%.

The low initial and declining share of the leading sector in total expenditure means that 40 years of 6%-per-year productivity growth in the leading sector has only a very limited impact on the total economy. After 40 years, the total productivity in the economy as a whole is only 2.54% higher than if the leading sector had not existed at all. Rapid productivity growth in the leading sector has next to no effect on productivity growth in the economy as a whole, because the salience of the leading sector falls, and the salience of other sectors resistant to productivity improvement rises, as technology advances. This is Baumol and Bowen's (1966) "cost disease" scenario: innovations become less and less important because the innovation-resistant share of the economy rises over time. Indeed, as time passes the rate of aggregate growth converges to the rate of growth in the productivity-resistant rest of the economy.

By contrast, with a price elasticity of 4 the expenditure share of the leading sectors grows rapidly from its original value of 2%. With a productivity-growth wedge of 5% per year, the leading-sector share of spending surpasses 10% by year 12 and 30% by year 20, and reaches 89% by year 40. As the spending share of the leading sectors rises, aggregate productivity growth rises too: from 1.1% per year at the start to 1.4% per year by year 10, 2.4% per year by year 20, 4.2% per year by year 30, and 5.4% per year by year 40. The impact on the aggregate economy is enormous: the TFP after 40 years is 113% higher than it would have been had the leading sector never existed.

There is only one reason for the sharp difference in the effects of innovation in the leading sector: the different price elasticities of demand for leading-sector products in the two scenarios. The initial shares of leading-sector products in demand, the rate of technology improvement in the leading sector, and the duration of the technology boom are all the same. But when the demand for leading-sector products is price-elastic, each advance in technology and reduction in the leading sector's costs raises the salience of the leading sector in the economy and thus brings the proportional rate of growth of the aggregate economy closer to the rate of growth in the leading sector itself. By the end of the 40-year period of these simulations, the scenario with the price elasticity of 4 has seen the leading sectors practically take over the economy, and dominate demand. This is the "true economic revolution" scenario: not only does productivity growth accelerate substantially and material welfare increase, but the structure of the economy is transformed as the bulk of the labor force shifts into producing leading-sector products and the bulk of final demand shifts into consuming leading-sector products.

What determines whether demand for a leading sector's products is price-inelastic (in which case we are in Baumol and Bowen's "cost disease" scenario in which technological progress in the leading sector barely affects the aggregate economy at all) or price-elastic (in which case we are in the "economic revolution" scenario, and everything is transformed)? What determines the income and price elasticities of demand for the high-tech goods that are the products of our current leading sectors?

4.2 HOW USEFUL WILL COMPUTERS BE?

What factors determine what the ultimate impact of these technologies will be? What is there that could interrupt a bright forecast for productivity growth over the next decade? There are three possibilities: The first is the end of the era of technological revolution—the end of the era of declining prices of IT capital. The second is a steep fall in the share of total nominal expenditure devoted to IT capital. And the third is a steep fall in the social marginal product of investment in IT—or, rather, a fall in the product of the social return on investment and the capital–output ratio. The important thing to focus on in forecasting is that none of these have happened: In 1991–1995 semiconductor production was 0.5% of nonfarm business output; in 1996–2000 it averaged 0.9%. Nominal spending on IT capital rose from about 1% of GDP in 1960 to about 2% by 1980 to about 3% by 1990 to between 5% and 6% by 2000. Computer and semiconductor prices declined at 15–20% per year from 1991 to 1995 and at 25–35% per year from 1996 to 2000.

However, whether nominal expenditure shares will continue to rise in the end hinges on how useful data processing and data communication products turn out to be. What will be the elasticity of demand for high-technology goods as their prices continue to drop? The greater is the number of different uses found for high-tech products as their prices decline, the larger will be the income and price elasticities of demand—and thus the stronger will be the forces pushing the expenditure share up, not down, as technological advance continues. All of the history of the electronics sector suggests that these elasticities are high, nor low. Each successive generation of falling prices appears to produce new uses for computers and communications equipment at an astonishing rate.

The first, very expensive computers were seen as good at performing complicated and lengthy sets of arithmetic operations. The first leading-edge applications of large-scale electronic computing power were military: the burst of innovation during World War II that produced the first one-of-a-kind hand-tooled electronic computers was totally funded by the war effort. The coming of the Korean War won IBM its first contract to actually deliver a computer: the million-dollar Defense Calculator. The military demand in the 1950s and the 1960s by projects such as Whirlwind and SAGE (Semi-Automatic Ground Environment)—a strategic air defense system—both filled the assembly lines of computer manufacturers and trained the generation of engineers that designed and built them.

The first leading-edge civilian economic applications of large—for the time (the 1950s)—amounts of computer power came from government agencies like the Census and from industries like insurance and finance, which performed lengthy sets of calculations as they processed large amounts of paper. The first UNIVAC computer was bought by the Census Bureau. The second and third orders came from A.C. Nielson Market Research and the Prudential Insurance Company. This second, slightly cheaper generation of computers was used not to make sophisticated calculations, but to make the extremely simple calculations needed by the Census and by the human-resource departments of large corporations. The Census Bureau used computers to replace their electromechanical tabulating machines. Businesses used computers to do the payroll, report-generating, and record-analyzing tasks that their own electromechanical calculators had previously performed.

The still next generation of computers—exemplified by the IBM 360 series—were used to stuff data into and pull data out of databases in real time—airline reservation processing systems, insurance systems, inventory control. It became clear that the computer was good for much more than performing repetitive calculations at high speed. The computer was much more than a calculator, however large and however fast. It was also an organizer. American Airlines used computers to create its SABRE automated reservations system, which cost as much as a dozen airplanes. The insurance industry automated its back-office sorting and classifying.

Subsequent uses have included computer-aided product design, applied to everything from airplanes designed without wind tunnels to pharmaceuticals designed at the molecular level for particular applications. In this area and in other applications, the major function of the computer is not as a calculator, a tabulator, or a database manager, but as a what-if machine. The computer creates models of what would happen if the airplane, the molecule, the business, or the document were to be built up in a particular way. It thus enables an amount and a degree of experimentation in the virtual world that would be prohibitively expensive in resources and time in the real world.

The value of this use as a what-if machine took most computer scientists and computer manufacturers by surprise. None of the engineers designing software for the IBM 360 series, none of the parents of Berkeley UNIX, nobody before Dan Bricklin programmed Visicalc had any idea of the utility of a spreadsheet program. Yet the invention of the spreadsheet marked the spread of computers into the office as a what-if machine. Indeed, the computerization of Americas white-collar offices in the 1980s was largely driven by the spreadsheet program's utility—first Visicalc, then Lotus 1-2-3, and finally Microsoft Excel.

For one example of the importance of a computer as a what-if machine, consider that today's complex designs for new semiconductors would be simply impossible without automated design tools. The process has come full circle. Progress in computing depends upon Moore's law; and the progress in semiconductors that makes possible the continued march of Moore's law depends upon progress in computers and software.

As increasing computer power has enabled their use in real-time control, the domain has expanded further as lead users have figured out new applications. Production and distribution processes have been and are being transformed. Moreover, it is not just robotic auto painting or assembly that have become possible, but scanner-based retail quick-turn supply chains and robot-guided hip surgery as well.

In the most recent years the evolution of the computer and its uses has continued. It has branched along two quite different paths. First, computers have burrowed inside conventional products as they have become embedded systems. Second, computers have connected outside to create what we call the World Wide Web: a distributed global database of information all accessible through the single global network. Paralleling the revolution in data-processing capacity has been a similar revolution in data communications capacity. There is no sign that the domain of potential uses has been exhausted.

One would have to be pessimistic indeed to forecast that all these trends are about to come to an end. One way to put it is that modern semiconductor-based electronics technologies fit Bresnahan and Trajtenberg's (1995) definition of a *general-purpose technology*—one useful not just for one narrow class but for an extremely wide variety of production processes, one for which each decline in price appears to bring forth new uses, one that can spark off a long-lasting major economic transformation. There is room for computerization to grow on the intensive margin, as computer use saturates potential markets like office work and email. But there is also room to grow on the extensive margin, as microprocessors are used for tasks (like controlling hotel-room doors or changing the burn mix of a household furnace) that few, two decades ago, would have thought of.

5. Additional Considerations

Moreover, the analysis so far has left out a substantial number of important considerations.

5.1 PREVIOUS INDUSTRIAL REVOLUTIONS

The first of these is that previous industrial revolutions driven by general-purpose technologies have seen an initial wave of adoption followed by rapid TFP growth in industries that use these new technologies as businesses and workers learn by using. So far this has not been true of our current wave of growth. As Robert Gordon (2002) has pointed out at every opportunity, there has been little if any acceleration of TFP growth outside of the making of high-tech equipment itself: the boosts to labor productivity look very much like what one would expect from capital deepening alone, not what one would expect from the fact that the new forms of capital allow more efficient organizations.

Paul David (1991) at least has argued that a very large chunk of the long-run impact of technological revolutions does emerge only when people have a chance to thoroughly learn the characteristics of the new technology and to reconfigure economic activity to take advantage of it. In David's view, it took nearly half a century before the American economy had acquired enough experience with electric motors to begin to use them to their full potential. By his reckoning, we today are only halfway through the process of economic learning needed for us to even begin to envision what computers will be truly useful for.

Moreover, as Crafts (2002) argues, the striking thing is not that there was a "Solow paradox" of slow productivity growth associated with computerization, but that people did not expect the economic impact to start slow and gather force over time. As he writes, "in the early phases of general purpose technologies their impact on growth is modest." It has to be modest: "the new varieties of capital have only a small weight relative to the economy as a whole." But if they are truly general-purpose technologies, their weight will grow.

Susanto Basu's comment on this paper suggests that we are finally beginning to see Paul David's point begin to have force. As time passes, it looks like a larger and larger share of the TFP growth acceleration is coming in industries outside of the high-tech sector—in industries that are learning how to use IT products to boost their own efficiency of operations.

5.2 ADJUSTMENT COSTS

Basu, Fernald, and Shapiro (2001) estimate that because of adjustment costs productivity growth in the second half of the 1990s *undershot* the long-run technology trend by half a percentage point per year or more. Our standard models tell us that investment is more or less stable over time because adjustment costs are substantial: to invest 10% of national product in equipment this year and 2% the next is much worse than investing a steady 6% in equipment. But the 1990s saw sudden, unprecedented, large shifts in real investment shares. If our standard explanations of why investment does not swing more wildly are correct, then the penalties enforced by adjustment costs on American economic growth in the late 1990s must have been relatively large.

As Martin Baily (2002) has observed, there is independent evidence for these adjustment costs: "microeconomic analyses of plants and firms find substantial adjustment costs to investment and lags between investment and productivity." Thus it is highly naive to follow "the growth accounting approach," and to assume that "increases in capital intensity have an impact on productivity in the same year" or even the same five-year period in which they occur.

5.3 ECONOMIC GOVERNANCE

There is, however, the one remaining hole card: the pessimistic one. The government's role in a market economy is to provide the underlying definitions of property rights, mechanisms of rights enforcement, and corrections of externalities so that the price signals sent by the market to firms correspond to economic efficiency and social values. As the structure of the economy changes, surely the proper government-provided institutional underpinnings must change too. But is government able to fulfill this market-structuring task?

The macroeconomist tends to foresee a future of falling high-tech prices, rising expenditure shares, rapidly growing capital–output ratios, and fast labor productivity growth. Yet as one looks at IT, one cannot help but be struck by the fact that the most far-reaching and important consequences may well be microeconomic. Issues of the benefits from the extent of the market, of price discrimination and the distribution of economic well-being, of monopoly, and of the interaction of intellectual property with scientific communication and research are all very important and very complicated. And if governments fail to properly structure the micro marketplace, then optimistic macro conclusions will be immediately cast into doubt.

It is obvious that the creation of knowledge is a cumulative enterprise: Isaac Newton said that the only reason he was able to see farther than others was that he stood on the shoulders of giants. Whenever we consider the importance of property rights over ideas in giving companies incentives to fund research and development, we need to also consider the importance of free information exchange and use in giving researchers the power to do their jobs effectively. Can governments construct intellectual-property systems that will both enhance information exchange and provide sufficient monetary incentives? It is an open question.

One possible solution may be price discrimination. In the past, price discrimination—charging one price to one consumer and a different price for essentially the same good to another consumer—has been seen as a way for monopolies to further increase their monopoly profits. In the information age the background assumption may be different. We may come to see price discrimination as an essential mechanism for attaining economic efficiency and social welfare.

Third, if we call the economy of the past two centuries primarily *Smithian*, the economy of the future is likely to be primarily *Schumpeterian*. In a Smithian economy, the decentralized market does a magnificent job (if the initial distribution of wealth is satisfactory) at producing economic welfare. Since goods are *rival*—my sleeping in this hotel bed tonight keeps you from doing so—one person's use or consumption imposes a social cost: since good economic systems align the incentives facing individuals with the effects of their actions on social welfare, it makes sense to distribute goods by charging prices equal to marginal social cost. Since goods are *excludable*—we have social institutions to enforce property rights, in the case of my hotel room the management, the police, and the courts—it is easy to decentralize decision making and control, pushing responsibility for allocation away from the center and to the more entrepreneurial periphery where information about the situation on the ground is likely to be much better.

In a Schumpeterian economy, the decentralized market does a much less good job. Goods are produced under conditions of substantial increasing returns to scale. This means that competitive equilibrium is not a likely outcome: the canonical situation is more likely to be one of natural monopoly. But natural monopoly does not meet the most basic condition for economic efficiency: that price equals marginal cost. However, prices cannot be *forced* to be equal to marginal cost, because then the fixed setup costs are not covered. Relying on government subsidies to cover fixed setup costs raises problems of its own: it destroys the entrepreneurial energy of the market and replaces it with the group-think and red-tape defects of admininstrative bureaucracy. Moreover, in a Schumpeterian economy it is innovation that is the principal source of wealth—and temporary monopoly power and profits are the reward needed to spur private enterprise to engage in such innovation. The right way to think about this complex set of issues is not clear. The competitive paradigm cannot be fully appropriate. But it is not clear what is.

Consider, for example, the U.S. Gilded Age toward the end of the nineteenth century. The Gilded Age saw the coming of mass production, the large corporation, the continent-wide market, and electric power to the United States. You needed more than the improvements in production technology that made possible the large-scale factory in order to arrive at the large industrial organization and the high-productivity, mass-production economy. From our viewpoint today we can look back and say that in the United States this economic transformation rested on five things:

Limited liability.
The stock market.
Investment banking.
The continent-wide market.
The existence of an antitrust policy.

Legal and institutional changes—limited liability, the stock market, and an investment banking industry—were needed to assemble the capital to build factories on the scale needed to serve a continental market. Without limited liability, individual investors would have been unwilling to risk potentially unlimited losses from the actions of managers they did not know and could not control. Without the stock and bond markets, investors would have been less willing to invest in large corporations because of the resulting loss of liquidity. Without investment banking, investors' problem of sorting worthwhile enterprises from others would have been much more difficult.

Moreover, political changes—the rise of antitrust—were needed for two reasons. The first was to try to make sure that the enormous economies of scale within the grasp of the large corporation were not achieved at the price of replacing competition by monopoly. The second was the political function of reassuring voters that the growing large corporations would be the economy's servants rather than the voters' masters.

Last, institutional changes were needed to make sure that the new corporations could serve a continental market. For example, think of Swift Meatpacking. Swift's business was based on a very good idea: mass-slaughter the beef in Chicago, ship it dressed to Boston, and undercut local small-scale Boston-area slaughterhouses by a third at the butcher shop. This was a very good business plan. It promised to produce large profits for entrepreneurs and investors and a much better diet at lower cost for consumers. But what if the Massachusetts legislature were to require for reasons of health and safety that all meat sold in Massachusetts be inspected live and on the hoof by a Massachusetts meat inspector in Massachusetts immediately before slaughter?

Without the right system of governance—in this case U.S. federal preemption of state health and safety regulation affecting interstate commerce—you wouldn't have had America's Chicago meatpacking industry (or Upton Sinclair's *The Jungle*). That piece of late-nineteenth century industrialization wouldn't have fallen into place.

Because American institutions changed to support, nurture, and manage the coming of mass production and the large-scale business enterprise chronicled by Alfred Chandler—and because European institutions by and large did not—it was America that was on the cutting edge at the start of the twentieth century. It was America that was "the furnace where the future was being forged," as Leon Trotsky once said.

What changes in the government-constructed underpinnings of the market economy are needed for it to flourish as the economic changes produced by computers take hold? Optimistic views of future macro productivity growth assume that government will—somehow—get these important micro questions right.

6. *Conclusion*

The main argument of this paper has been that standard growth models predict a bright future. Increased TFP in the IT capital-goods-producing sector, coupled with extraordinary real capital deepening as the quantity of real investment in IT capital bought by a dollar of nominal savings grows, has driven the productivity growth acceleration of the later 1990s, and promise to drive equal or faster growth in the next decade. The extraordinary pace of invention and innovation in the IT sector has generated real price declines of between 10% and 20% per year in information processing and communications equipment for nearly forty years so far. These extraordinary cost declines have made a unit of real investment in computer or communications equipment absurdly cheap, and hence made the quantity of real investment and thus capital deepening in IT capital absurdly large.

Continued declines in the prices of IT capital mean that a constant nominal flow of savings channeled to such investments brings more and more real investment. The social return to IT investment would have to drop suddenly to nearly zero, the share of nominal investment spending devoted to IT capital would have to collapse, or technological progress in IT would have to slow drastically, for labor productivity growth in the next decade to reverse itself and return to its late 1970s or 1980s levels. Yet there are no technological reasons for the pace of productivity increase in our economy's leading sectors to decline over the next decade or so.

Moreover, the future may well be brighter than the simplest models suggest. First, elasticity of demand for IT goods is likely to remain high. A high elasticity of demand for IT goods means that as prices fall, expenditure shares will not remain constant but will rise, boosting growth. Second, the long-time trend shows the share of national income attributable to the returns on the existing IT capital stock rising as well. This rise should boost growth as well. Third, Basu, Fernald, and Shapiro (2001) take adjustment costs in investment seriously, and conclude that late-1990s growth undershot trend growth by about 0.5 percentage points per year. Fourth, David (1991) has argued for decades that general-purpose technologies boost labor productivity in successive stages, with the largest boost coming after the technology has diffused throughout the economy.

Is there any reason to be pessimistic? Only a fear of large-scale governmental failure in setting forth the institutional framework to support information-age markets could lead to a pessimistic forecast of future growth.

REFERENCES

Baily, M. (2002). The new economy: Post-mortem or second wind? *Journal of Economic Perspectives,* 16(2) (Spring):3–22.

Basu, S., J. Fernald, and M. Shapiro. (2001). Productivity growth in the 1990s: Technology, utilization, or adjustment? Cambridge, MA: National Bureau of Economic Research. NBER Working Paper 8359.

Baumol, W., and W. Bowen. (1966). Performing Arts: The Economic Dilemma. New York: Twentieth Century Fund.

Bresnahan, T., and M. Trajtenberg. (1995). General purpose technologies: "Engines of growth"? *Journal of Econometrics* 65:83–108.

Cohen, S., J. B. DeLong, and J. Zysman. (2000). Tools for thought: What is new and important about the e-conomy. Berkeley: BRIE. Working Paper 138.

Crafts, N. F. R. (1985). *British Economic Growth during the Industrial Revolution.* Oxford: Oxford University Press.

———. (2000). The Solow productivity paradox in historical perspective. London: London School of Economics. Mimeo.

David, P. (1991). Computer and dynamo: The modern productivity paradox in a not-too-distant mirror. In *Technology and Productivity: The Challenge for Economic Policy.* Paris: OECD.

DeLong, J. B. (1991). Liquidation cycles and the Great Depression. Cambridge, MA: Harvard University. Mimeo.

Fischer, S. (1988). Symposium on the slowdown in productivity growth. *Journal of Economic Perspectives* 2(4, Fall):3–7.

Gordon, R. (2000a). Does the "new economy" measure up to the great inventions of the past? *Journal of Economic Perspectives* 14(4):49–74.

———. (2000b). Interpreting the "one big wave" in U.S. long term productivity growth. Cambridge, MA: National Bureau of Economic Research. NBER Working Paper 7752.

———. (2002). Technology and economic performance in the American economy. Cambridge, MA: National Bureau of Economic Research. NBER Working Paper 8771.

Greenwood, J., and B. Jovanovic. (1999). The information technology revolution and the stock market. *American Economic Review* 89(2, May):116–122.

Griliches, Z. (1988). Productivity puzzles and R & D: Another nonexplanation. *Journal of Economic Perspectives* 2(4, Fall):9–21.

Jorgenson, D. (1988). Productivity and postwar U.S. economic growth. *Journal of Economic Perspectives* 2(4, Fall):23–41.

———, M. Ho, and K. Stiroh. (2001). Projecting productivity growth: Lessons from the U.S. growth resurgence. Cambridge, MA: Harvard University. Mimeo.

Kosterlitz, J. (2002). What went wrong? *National Journal,* January 4.

Krugman, P. (1994). *The Age of Diminished Expectations*, 4th ed. Cambridge, MA: The MIT Press.
Landes, D. (1969). *The Unbound Prometheus*. Cambridge: Cambridge University Press.
McKinsey Global Institute. (2001). *U.S. Productivity Growth, 1995–2000*. Washington: McKinsey.
Oliner, S., and D. Sichel. (2000). The Resurgence of Growth in the Late 1990s: Is Information Technology the Story? *Journal of Economic Perspectives* 14(4):3–22.
Staiger, D., J. H. Stock, and M. W. Watson. (1997). The NAIRU, unemployment and monetary policy. *Journal of Economic Perspectives* 11(1, Winter):33–49.
Stiroh, K. (2001). Information technology and the U.S. productivity revival: What do the industry data say? New York: Federal Reserve Board of New York. Working Paper 115.
Walton, S. (1992). *Made in America: My Story*. New York: Bantam Books.
Whelan, K. (2000a). A guide to the use of chain-aggregated NIPA data. Washington: Federal Reserve Board. Mimeo.
———. (2001). A two sector approach to modeling NIPA data. Washington: Federal Reserve Board. Mimeo.

Comment

SUSANTO BASU
University of Michigan and NBER

Economists are not particularly good at understanding large, one-time structural changes. We are not notably successful at working in real time. And despite much effort, especially over the past 20 years, we still have no generally accepted understanding of the underlying sources of technical change. Any single paper that tries to tackle all three problems at once is both incredibly ambitious, and almost surely destined to failure. Thus, it is not surprising that DeLong's paper does not succeed in explaining the enormously important and extremely puzzling behavior of productivity growth in the United State since the mid-1990s—no single paper could be expected to do so. But it illuminates and elucidates most of the important issues that need to be addressed if we are to succeed as a profession in explaining this very important puzzle. In keeping with the spirit of the paper, my comment tries to put on the table one more fact that needs to be explained if we are to have a tolerably complete understanding of the recent productivity acceleration.

Before doing so, however, I want to shift the discussion from labor productivity, which is the focus of DeLong's paper, to technical change (which, under standard conditions, is equivalent to total factor productivity). The reason is twofold. First, a focus on labor productivity conflates impulses and propagation mechanisms, and gives the somewhat misleading impression that we understand the sources of productivity growth better than we actually do. DeLong's paper actually takes the rate of technical change and its recent acceleration as given, and combines these with various standard models of investment. But the fact that we understand how falling capital-goods prices can induce capital accumulation does not mean that we truly understand the reasons for higher labor or total factor productivity growth in the United States. Second, focusing on technical change helps put an important fact into sharper focus.[1]

To understand the claims I make below, it's necessary to do a bit of growth accounting. Assume that the economy comprises just two sectors, one producing information technology (IT)—here identified with the production of computers, semiconductors, and telecommunications equipment—and the other producing everything else (called NT, for non-IT). Furthermore, assume that production in each sector can be represented using a production function for gross output (Q), using capital (K), labor (L), intermediate inputs (M), technology (T), and gross investment (I):

$$Q_j = F(K_j, L_j, M_j, T_j, I_j), \qquad j = \text{IT, NT}. \tag{1}$$

The only variable that is not completely standard is investment. Its presence is meant to capture the costs of adjusting the capital stock, here taken to be internal (output-reducing) rather than external, as is often assumed. One should think of these "costs" in a very broad sense as including complementary (but otherwise unmeasured) investments. Thus, rather than workers and machines being idled as new capital is installed, perhaps machines are idled and workers sent to be retrained on the new equipment. Furthermore, the retraining—both formal and on the job—can last far longer and be much more costly than a quick office or factory shutdown.[2]

Letting lowercase letters represent natural logs of their uppercase counterparts, letting starred variables be their steady-state values, and approximating time derivatives by finite differences, equation (1) implies that the unobserved technical change in each sector can be approximated as[3]

1. Sometimes the use of labor productivity is justified by claiming that it gives a way to take account of unmeasured output in the household sector. But Basu and Fernald (2001) argue that total factor productivity is in fact a superior welfare measure as well: by subtracting the change in labor input from output growth, the TFP calculation implicitly values home production at its marginal opportunity cost in forgone market wages.
2. In principle, one can add costs of adjusting all the inputs, especially labor. However, empirically the costs of adjusting capital seem much larger, and the object of the exercise is to understand the late 1990s, which were characterized by extremely high rates of capital investment.
3. See Basu, Fernald, and Shapiro (2001) for a derivation.

$$\Delta t_j = \Delta q_j - \mu_j(S_j^K \, \Delta k_j + s_j^L \, \Delta l_j + s_j^M \, \Delta m_j) - \frac{F_5^* I^*}{Q^*} \Delta i_j, \tag{2}$$

where the s_j terms are the shares of input costs in the total revenue of sector j, and μ_j is the ratio of price to marginal cost in that sector.

From the discussion above, we expect that $F_5 < 0$. Hence, the last term in equation (2) is actually positive as long as investment is growing. The intuition is that the process of installing capital is in fact an output of the firm, although it is not recorded as such because it is not a market good. If the firm installs capital instead of producing market output then measured output falls, but total output does not fall as much and may actually rise. The last term in equation (2) adds back this unmeasured output growth, which means that in times of high investment *growth* technology is actually rising faster than the usual calculation would lead us to believe.[4]

One can aggregate technical change in the two sectors to get a measure of overall technology growth using Domar (1961) aggregation, extended to allow for imperfect competition:

$$\Delta t = \sum_{j=\text{IT,NT}} \frac{P_j Y_j}{P_{\text{IT}} Y_{\text{IT}} + P_{\text{NT}} Y_{\text{NT}}} \frac{\Delta t_j}{1 - \mu_j s_j^M}, \tag{3}$$

where $P_j Y_j$ is the total revenue in sector j.

Basu, Fernald, and Shapiro (2001) implement this framework, with additional controls for variable factor utilization. Using their industry-level results (not reported in their paper), one can compute time series of technical change for the aggregate and the contributions of IT and NT to that aggregate change. Following the recent convention (e.g., Stiroh, forthcoming), I take the IT-producing industries to be SIC 35 and 36. Table 1 shows the means of these series for various sample periods.

Lines 2 and 3 show the results for the first and second halves of the 1990s. In both subperiods IT has an importance disproportionate to its size (about 5.5% of the economy). However, since the IT-producing sectors are such a small share of the economy, most of the growth in technical change for the economy as a whole is driven by technical change in the non-IT-producing sectors. Subtracting line 2 from line 3 gives the results for the late 1990s acceleration on line 4. The surprising result is that of the 1.9-percentage-point acceleration in technical change, only 0.3 percentage points came from IT. The acceleration of technical change in the economy

4. This idea is similar to the "1974" hypothesis of Yorukoglu and Greenwood (1997). In terms of implementation, the main difference is their suggestion that the costs of adjustment may last for many years, so one should have a number of lags of investment as well as current investment in equations (1) and (2).

Table 1 TECHNICAL CHANGE BY INDUSTRY GROUP AND SUBPERIOD

	Δt		
Subperiod	*Total*	*From IT*	*From NT*
1987–1999	2.0	0.5	1.5
1990–1995	1.2	0.3	0.9
1995–1999	3.1	0.7	2.4
Late 1990s acceleration	1.9	0.3	1.6

in the second half of the 1990s was due to something other than faster technical change in the sectors producing computers and telecommunications equipment.

This result, seemingly intuitive, is actually quite difficult to reconcile with the usual formalization of production. DeLong cites Stiroh (forthcoming) as saying that some of the labor productivity gains in the economy come from IT-using sectors. But the same should not be true of total factor productivity. According to the usual model of production—e.g., the production function for NT in equation (1)—more or better inputs are movements *along* the production function, not shifts *of* the production function. As the usual dictum has it, "cheaper inputs don't shift production functions."

There are at least four stories that might explain this observation. Two attribute it to mismeasurement, and two to actual technical change. The first mismeasurement story is that IT capital quality is growing faster than measured in the statistics, so there is effectively more IT investment than is actually measured, with the unmeasured investment wrongly showing up as faster technical change. This story seems unlikely, because a number of papers by Erik Brynjolfsson and his coauthors (e.g., Brynjolfsson, Hitt, and Yang, 2000) show that the increased investment raises productivity five or more years after the investment is made, not immediately as one would expect with straightforward mismeasurement. The second story is actually the preferred explanation of Brynjolfsson, Hitt, and Yang (2000) for their results. They attribute the increases in productivity and stock market valuation after IT investment to unmeasured complementary investments—for example, training, learning-by-doing, and general research on using IT effectively. Here the question is whether these unmeasured investments are nonrival. If they are, at least in part—and many of them sound like some form of knowledge—then they do constitute technical change in the sense that other firms can boost production without incurring similar resource costs. Of course innovations may take some time to spread to other firms—this temporary competitive advantage is presumably why the private investment is worthwhile—but slow diffusion of technology is actually a reason to be hopeful, for it indicates that many of the social returns to these complementary investments may lie ahead.

The two interpretations of the Table 1 results as denoting actual technological change are quite similar to the nonrival complementary investment story. The first takes seriously the idea of computers and telecommunications as a new general-purpose technology (GPT) (see Helpman and Trajtenberg, 1998). As these models assume and economic history shows clearly, a GPT needs to be combined with various specific innovations in order to yield large production efficiencies. These innovations often require substantial time, and secondary innovations of their own, to change the structure of production. The second interpretation—really a special case of the GPT story—is that there is capital-biased (Solow-neutral) technical change. In principle the derivative of equation (1) with respect to T may depend on all the other variables, including K. Jorgenson (1988) has used the idea of energy-biased technical change to explain the productivity slowdown that started in the late 1960s or early 1970s. One problem with this story has always been the absence of a major speedup in productivity growth following the collapse of oil prices in the mid-1980s. A simpler, unified explanation is the hypothesis of directed technical change: the idea that innovation is directed towards economizing on expensive inputs while being profligate with cheap ones. In the 1950s and 1960s the cheap input was oil; in the 1990s it was IT hardware. But in any given subperiod, directed technological change may be indistinguishable from factor-biased technical progress. It is in this sense that the second explanation is a special case of the first.

Thus, my intuition is that the key to understanding the productivity acceleration of the 1990s is making progress on modeling the process of directed technological change in the presence of a new GPT and, especially, confronting the models with detailed industry- and firm-level data. There are now some excellent overviews on how one might go about these difficult tasks—see, especially, Bresnahan (2001). But it is fair to say that these appealing stories have not had many empirical successes. If the stories are right, then technical change in the non-IT sectors have been driven by technical progress in IT, which has led to ever-falling prices for IT equipment. I tried to estimate the reduced-form model, correlating the industry residuals in the non-IT sector on lagged technical change in IT using the industry-level data used to construct Table 1. I found no evidence of a strong or stable relationship between the two. One can come up with many reasons why this relatively aggregative approach might not be successful. But the failure is disappointing nonetheless, because it means that there is little direct evidence for some of the most appealing economic models that might explain the major acceleration in productivity in the mid-1990s. Despite DeLong's nice paper, this area still abounds with questions awaiting answers.

REFERENCES

Basu, S., and J. G. Fernald. (2001). Aggregate productivity and aggregate technology. *European Economic Review* 46(6):963–991.

Basu, S., J. G. Fernald, and M. D. Shapiro. (2001). Productivity growth in the 1990s: Technology, utilization, or adjustment? *Carnegie-Rochester Conference Series on Public Policy* 55(1, December):117–165.

Bresnahan, T. (2001). The mechanisms of information technology's contribution to economic growth. Manuscript.

Brynjolfsson, E., L. M. Hitt, and S. Yang. (2000). Intangible assets: How the interaction of computers and organizational structure affects stock market valuations. MIT, working paper.

Domar, E. D. (1961). On the measurement of technical change. *Economic Journal* 71(December):710–729.

Helpman, E., and M. Trajtenberg. (1998). Diffusion of general purpose technologies. In *General Purpose Technologies and Economic Growth*, E. Helpman (ed.). Cambridge, MA: The MIT Press.

Jorgenson, D. W. (1988). Productivity and postwar U.S. economic growth. *Journal of Economic Perspectives* 2(4, Fall):23–41.

Stiroh, K. J. (forthcoming). Information technology and the U.S. productivity revival: What do the industry data say? *American Economic Review.*

Yorukoglu, M., and J. Greenwood. (1997). 1974. *Carnegie-Rochester Conference Series on Public Policy* 46(June):49–95.

Comment

BOYAN JOVANOVIC
New York University

1. Introduction

Bradford DeLong offers us an entertaining, well-written, and balanced introduction to the "new economy." We all hope, of course, that in the next decade or two we will see the economy return to the high growth rates that it experienced in the late 1990s. As consumers and workers, we hope that the productivity slowdown of the 1970s was a symptom of an investment episode that will pay off in the near future in the form of cheaper new products and higher wages for our labor services. And as a growth theorist, I root for the new economy because Schumpeter-style creative destruction is, to me, much more exciting to think about than an economy in a steady forward creep.

Brad's assessment is balanced: On the one hand, the productivity slowdown of the 1970s is still unexplained, and it is not at all clear that the slowdown represented a conscious sacrifice of output in return for expected future growth. On the other hand, Brad says that the rapid productivity growth of the late 1990s will probably persist for some time because

1. producers will rely even more on IT capital than they do at present (Section 4), and
2. Moore's Law will probably continue to hold for a decade or longer (Sections 2 and 4).

I agree with almost everything that Brad says in this paper, and I shall merely expand on these two specific points.

2. *The Growing Share of IT Capital*

As Brad observes in Section 5 of his paper, the GPTs of the past should provide us with a clue about what will happen in the coming decades. Let us take a look at the electrification of the United States after 1890. Warren Devine (1983) shows that electricity made its impact on industry in two waves. The first wave did not use electricity to its fullest extent. It simply used electricity to drive the main shaft to which machinery was attached by belts. Only after 1910, when individual drive was introduced with each machine itself plugged into an electric socket, were the full benefits achieved.

Figure 1 plots the isoquants of two technologies. The first technology, *shafts,* was intensive in nonelectrical equipment; it relied on electricity only to power the main drive to which old machinery was attached exactly as it had been when the main drive was powered by a steam engine. As a source of energy per se, electricity was cheaper than steam, but the shafts technology attained only a fraction of the cost savings that electrification had to offer. To get the full potential, one had to use electrical equipment plugged into sockets. This *wires* technology allowed more freedom in the type of building used and in the layout of the equipment, and it avoided the noise, dirt, and clutter of a factory where the conduits of power were belts linking the machinery to the heavily greased shafts.

Assume that each technology is Leontieff and yields constant returns to scale. Initially electrical equipment is expensive. Let the initial isocost line be P_1–P_1. We choose the initial date so that shafts and wires are equally affordable, and points A and A^* both yield a level of output equal to Y. Since returns are constant, this indifference must hold at any budget line parallel to P_1–P_1.

Over time, the budget line twists anticlockwise to, e.g., P_1–P_2. There, more output can be produced at point C with wires than at point B with shafts (i.e., $Y_3 < Y_2$). Figure 2 leaves out the switching costs that held

Figure 1 THE SWITCH FROM SHAFTS TO WIRES

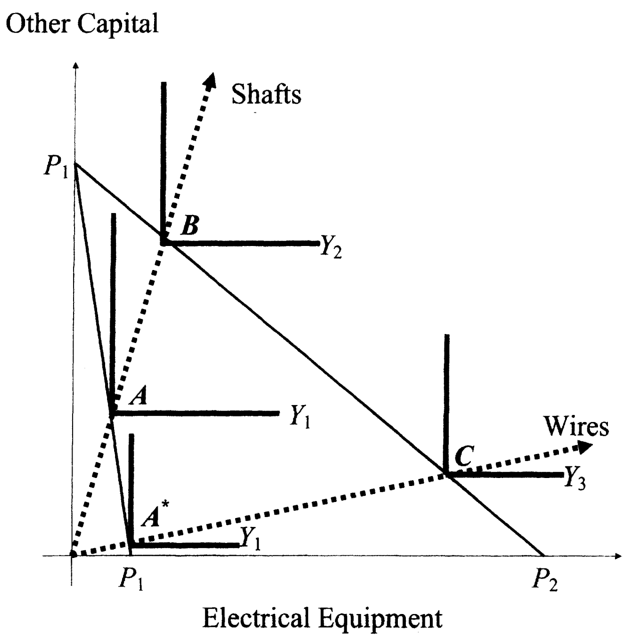

firms from using wires until the 1910–1930 period. As old and new factories were wired, productivity rose sharply.

Now fast-forward to the early 1970s when the U.S. economy began converting from the mainframe to the microcomputer. For "shafts" read "mainframe" and for "wires" read "PC," and the same story applies. The productivity growth of the late 1990s is, by most accounts, the result of this switch.

What does this say about today? Businesses shifted from shafts to wires between 1910 and 1930. They shifted from mainframes to personal computers after 1970. Each wave was followed by an episode of high productivity growth, and we hope that we are—or that we shall soon again be—in the middle of one such growth episode.

3. Computers vs. Electricity and Internal Combustion Engines

Moore's law translates into a rapid rate of price decline for computers and related equipment. Let us compare the price index for computers

Figure 2 PRICES AND CUMULATIVE QUANTITIES OF NEW-ECONOMY PRODUCTS

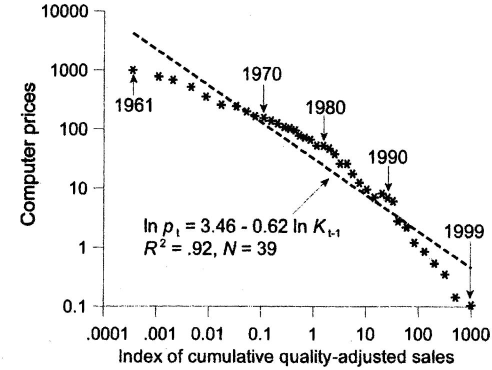

(a) Computer systems

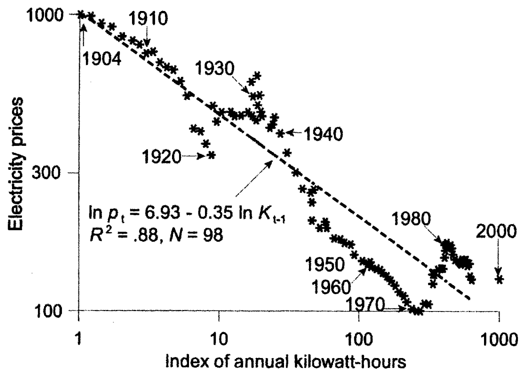

(b) Electricity

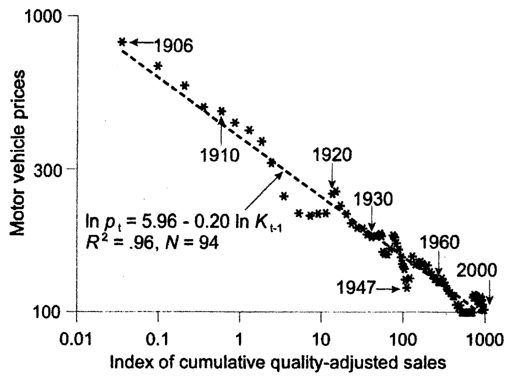

(c) Autos, trucks, and buses

with those of electricity and the internal combustion engine at a comparable stage in their development.

Gordon Moore stated his law in terms of time, but this does not tell us why computer prices fell. One hypothesis is that computer producers become more efficient as they sell more and more computers. Other GPTs also improved as they were applied. Management scientists often measure the application of a technology by the cumulative sales of its output. Let us therefore restate Moore's law in terms of cumulative-output growth instead of in terms of the passage of time.

Let p denote the price of a product. As the cost of producing it declines, so does its price, and all this is caused by the rise in the cumulative output, K, of all producers combined:

$$p = \left(\frac{K}{B}\right)^{-\beta}, \qquad (1)$$

where B is a constant. In logarithmic form, (1) reads

$$\ln p_t = \beta_0 - \beta \ln K_{t-1}, \qquad (2)$$

where $\beta_0 = -\beta \ln B$. Jovanovic and Rousseau (2002) estimate this equation for three GPTs: computers, electricity, and the internal combustion engine. Figure 2 presents pairwise combinations of $\ln p_t$ and $\ln K_{t-1}$ on an annual basis for each technology and plots a regression line through the points. The regression estimates are reported in Figure 2 itself.

The computer, displayed in panel (a), shows much the fastest decline in p, and the effect seems to be accelerating.[1] Panels (b) and (c)—electricity and cars—show a lower decline and no acceleration. Even if Moore's law were to slow to a quarter of its current pace, it still would dwarf the typical price decline in the other two GPTs. In fact, experts such as Meindl, Chen, and Davis (2001) tell us that Moore's law will hold at its current pace for at least another 20 years.

4. *Conclusion*

Moore's law is unique: Capital goods have never declined in price as fast as they are declining at present. On these grounds alone, productivity growth should for a while be well above its twentieth-century average.

Moreover, the use of IT involves a network externality in a way that

1. A description of how the price series were constructed is in Jovanovic and Rousseau (2002)—note 3 in the published version or note 2 in the NBER version.

neither electricity nor internal combustion did. This is apparent in the rising value of the Internet as more and more businesses get online. In spite of their low cost, computers and information systems have spread in full force only in a few rich countries. As people elsewhere join the network, they too will enjoy some of the gains that we saw in the United States in the 1990s, and the network effect will then lead to higher growth globally.

REFERENCES

Devine, W., Jr. (1983). From shafts to wires: Historical perspective on electrification. *Journal of Economic History* 43(2):347–372.

Jovanovic, B., and P. L. Rousseau. (2002). Moore's law and learning by doing. *Review of Economic Dynamics* 5(2, April):346–375. NBER Working Paper 8762. Cambridge, MA: National Bureau of Economic Research.

Meindl, J. D., Q. Chen, and J. A. Davis. (2001). Limits on silicon nanoelectronics for terascale integration. *Science* 293:2044–2049.

Discussion

Bob Gordon emphasized that even if Moore's law continues to hold, it is not certain what will happen to real computer investment, as this depends on the elasticity of demand for computing power. Gordon speculated that the key question is what "killer applications" on the demand side will be able to soak up the increase in computing power made possible by Moore's law and the fall in computer prices. He commented that the question of the extent to which productivity growth is driven by the old economy vs. the new economy is an interesting one, and remarked that several old-economy sectors such as landscaping are expanding in importance.

Brad DeLong remarked that the interaction of public policy, intellectual-property rights, and the development of new applications for computing power is a very interesting theme. He expressed surprise that the pricing and limitation of access to intellectual property is proceeding so slowly, especially considering the fact that some firms such as AOL–Time Warner are at one and the same time facilitating access and fighting to restrict access. However, Bob Hall disagreed with what he referred to as DeLong's pessimism on how the economy handles Schumpeterian goods. He remarked that while pessimism might be justified in theory, in practice the market for software functions quite successfully without government intervention.

Bob Hall commented further that, contrary to the assumption of the literature, there are two general-purpose technologies present rather than just one. He differentiated between the computer and the database as two distinct general-purpose technologies. He remarked that the success of companies such as WalMart, Dell, Southwest, and Ebay is built on the ability to keep track of huge amounts of data rather than on the computer as such. On Gordon's point with respect to new-economy- vs. old-economy-driven productivity growth, he was of the opinion that most sectors have not yet even begun to harness the power of modern IT. As an example, he compared WalMart with Kmart.

DeLong drew Hall's attention to the finding in the recent McKinsey U.S. productivity study that a large share of productivity growth at the micro level is due not to IT, but to competitive pressure from WalMart on the retail sector. He asked Hall how it is possible that this competitive pressure is not related to WalMart's use of IT. Hall responded that WalMart's success arises not just from its use of IT in the realm of logistics, but also from its success in human-resource management. Bob Gordon added that the advantage of Southwest Airlines over companies such as United lies in its employee relations rather than in its use of technology.

Mark Gertler asked DeLong to speculate further on the source of the productivity slowdown of the 1970s. He suggested a connection between low-frequency movements in unemployment and labor productivity as a starting point. DeLong replied that he saw the causality going from labor productivity to unemployment. He found plausible a story that linked low productivity growth to high unemployment through workers' expectations of wage growth that failed to keep in step with productivity.

Justin Wolfers invited Brad DeLong to speculate on the impact of IT on the progress of economic science. DeLong and Gordon both responded by referring to Griliches's comment that cheap computers meant too many regressions and too little thought put into specification. However, Bob Gordon remarked that the work of Stock and Watson would not be possible without the advances in computing power that have taken place.

Summing up, Brad DeLong agreed with Susanto Basu's discussion that it is embarrassing not to be able to say much about multifactor productivity growth. However, he commented that the danger in focusing on TFP growth is that capital deepening might be ignored. Looking forward, he speculated that a substantial amount of labor productivity growth will come about through capital deepening as a result of Moore's law and the falling price of IT. He also suggested that capital deepening might be responsible for the recent phenomenon of falling labor inputs at the same time as very little fall in real GDP. On the question of TFP growth outside the IT sector, he suggested that the computer is not merely a case of an input whose price is falling, but of an entirely new kind of input of the kind referred to by stories of the general-purpose-technology type. He agreed with both discussants that it makes sense to try to model multifactor productivity as generated by some form of human action. However, he noted that the importance of multifactor productivity suggests that if it were generated by some particular activity, economists would already have an idea of what that activity is, and how to encourage it.