

Productivity Growth in the 2000s

J. Bradford DeLong¹
University of California at Berkeley and NBER

March 2002

Draft 1.2

¹ I would like to thank the University of California at Berkeley for financial support, and Chad Jones, Dan Sichel, and Larry Summers for helpful discussions.

Abstract

The causes of the productivity growth slowdown of the 1970s remain mysterious. By contrast, nearly all agree that the cause of the productivity growth speed-up of the 1990s lie in the information technology sector. The extraordinary pace of invention and innovation in the information technology sector has generated real price declines of between ten and twenty percent per year for decades. Increased total factor productivity in the information technology capital goods-producing sector coupled with extraordinary real capital deepening as the quantity of real investment in information technology capital bought by a dollar of nominal savings grows have together driven the productivity growth acceleration of the later 1990s.

Will this new, higher level of productivity growth persist? The answer appears likely to be “yes.” The most standard of simple applicable growth models—that of Oliner and Sichel—predicts that the social return to information technology investment would have to suddenly and discontinuously drop to zero for the upward jump in productivity growth to reverse itself in the near future. More complicated models that focus in more detail on the determinants of investment spending or on the sources of increased total factor productivity appear to strengthen, not weaken, forecasts of productivity growth over the next decade.

I. Introduction

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 2.88 percent per year from 1947 to 1973. It averaged a growth rate of only 1.30 percent per year from 1973 to 1995. The deceleration in the growth rate of total real GDP was somewhat smaller: a matter of -1.18 percentage points per year in output rather than the -1.58 percentage points per year in labor productivity, as total real GDP growth slowed from 3.91 percent per year averaged over 1947:1 to 1973:2 to 2.73 percent per year averaged real over 1973:2 to 1995:1. 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. The productivity slowdown meant that, according to official statistics, Americans in 1995 were only 70 percent 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.²

In the second half of the 1990s American productivity picked itself up, and more-or-less resumed its pre-1973 pace. Between the beginning of 1995 and the semi-official NBER business cycle peak in March 2001, U.S. nonfarm-business output per person-hour worked grew at an annual rate of 2.80 percent per year. (Extending the sample through the 2001 recession to the likely trough point of 2001:4, the late-1990s growth rate is 2.69

² See Krugman (1994) for one interpretation of how the productivity slowdown made a big difference.

percent per year.) Between the beginning of 1995 and the semi-official NBER business cycle peak in March 2001, U.S. real GDP grew at a pace of 4.21 percent per year. (Extending the sample through the 2001 recession to the likely trough point of 2001:4, the late-1990s growth rate is 3.85 percent per year.) Non-economists tended to attribute a large chunk of fast late-1990s growth to “cyclical” factors,³ but economists 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.⁴ Moreover, as Susanto Basu, John Fernald, and Matthew Shapiro have powerfully argued, there are stronger reasons for thinking that the adjustment costs associated with moving to a more information technology 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.⁵ And the extremely rapid run-up of stock prices indicated that at least the marginal investor in equities anticipated that the acceleration of economic growth that started in the mid-1990s would last for decades or longer.⁶

The causes of the productivity slowdown of the 1973-1995 or so period remain disappointingly mysterious. Baily (2002) calls the growth-accounting literature on the slowdown “large but inconclusive.” No single factor provides a convincing and coherent

³ See, for example, Kosterlitz (2002).

⁴ See Gordon (2002) and Gordon (2000).

⁵ See Basu, Fernald, and Shapiro (2001).

⁶ See Greenwood and Jovanovic (1999).

explanation, and the residual position that a large number of growth-retarding factors suddenly happened to hit at once is unlikely.⁷

By contrast, nearly all agree on the causes of the productivity speed-up of 1995-2001: it is the result of the extraordinary wave of technological innovation in computer and communications equipment—solid-state electronics and photonics.⁸ 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 information technology—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... 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 [information technology capital goods] or use [information technology capital goods] most intensively, with no net contribution from other industries... relatively isolated from the [information technology] revolution.” Oliner and Sichel (2000) write that “the rapid capital deepening related to information technology capital accounted for nearly half of

⁷ See Fischer (1988), Griliches (1988), Jorgenson (1988), and Gordon (2000b) and (2002). Jorgenson (1988) convincingly demonstrates that the oil price shocks can 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) and (2002) finds such an explanation attractive.

⁸ The only major study taking a stand against this explanation is 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, they find next to no correlation between computer capital and labor productivity. When they weight industries by employment, they find a statistically significant and substantively important connection.

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 information technology capital-deepening and increased efficiency in the production of computing and communications 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.⁹

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 speed-up. This paper, however, will assume that this consensus is correct: that the productivity growth speed-up in the second half of the 1990s was the result of the technological revolutions in data processing and data communications. It will then go on to ask what the boom of the past seven years means for productivity growth in the next decade or so. Will the decade of the 2000s be more like the late 1990s, or more like the 1980s as far as growth in productivity and living standards is concerned?

The answer is that the smart way to bet is that the 2000s will be much more like the fast-growing late 1990s than like the 1980s. The extraordinary pace of invention and innovation in the information technology sector has generated real price declines of between ten and twenty percent per year in information processing and communications

⁹ However, Jorgenson, Ho, and Stiroh expect total real GDP growth to slow because of slower growth in hours worked—they forecast 1.1 percent per year growth in hours over the next decade, compared to 2.3 percent per year growth in hours from 1995 to 2000.

equipment for nearly forty years so far. There are no technological reasons for this pace of productivity increase in these leading sectors to decline over the next decade or so. In the consensus analysis, increased total factor productivity in the information technology capital goods-producing sector coupled with extraordinary real capital deepening as the quantity of real investment in information technology capital bought by a dollar of nominal savings grows have together driven the productivity growth acceleration of the later 1990s. 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 cost declines had made a unit of real investment in computer or communications equipment absurdly cheap, hence the quantity of real investment and thus capital deepening in information-technology capital absurdly large.

Thus the most standard of simple applicable growth accounting approaches predicts a bright future for the American economy over the next decade or so. Continued declines in the prices of information technology capital mean that a constant nominal flow of savings channeled to such investments will bring more and more real investment. As long as information technology capital earns the same rate of return as other capital, then labor productivity growth should continue very high. The social return to information technology investment would have to suddenly and discontinuously drop to nearly zero, or the share of nominal investment spending devoted to information technology capital would have to collapse, or both, for labor productivity growth in the next decade to reverse itself and return to its late 1970s or 1980s levels.

Moreover, additional considerations tend to strengthen, not weaken, forecasts of productivity growth over the next decade. It is very difficult to argue that the speculative excesses of the 1990s boom produced substantial upward distortions in the measured growth of potential output. The natural approach that economists use to model investment spending in detail—the approach used by Basu, Fernald, and Shapiro (2001)—tells us that times of rapid increase in real investment are times when “adjustment costs” are unusually high, and thus times when actual productivity growth undershoots the long-run sustainable trend. Both a look back at past economic revolutions driven by general-purpose technologies¹⁰ that were in their day analogous to the computer in their effects¹¹ and a more deeper look forward into the likely determinants of productivity growth suggest a bright future.

The pace of technological progress in the leading sectors driving the "new economy" is very rapid indeed, and will continue to be very rapid for the foreseeable future. Second, the computers, switches, cables, and programs that are the products of today's leading sectors are what Bresnahan and Trajtenberg (1985) call “general-purpose technologies,” hence demand for them is likely to be extremely elastic. Rapid technological progress brings rapidly falling prices. Rapidly falling prices in the contest of extremely elastic demand will produce rapidly-growing expenditure shares. And the economic salience of a leading sector--its contribution to productivity growth--is the product of the rate at which the cost of its output declines and the share of the products it makes in total demand.

¹⁰ See Bresnahan and Trajtenberg (1995).

Thus unless Moore's Law ceases to hold or the marginal usefulness of computers and communications equipment rapidly declines, the economic salience of the data processing and data communications sectors will not shrink, but grow.

This paper attempts to demonstrate these points in five sections, including this introduction. Section II lays out the aggregate macroeconomic facts about the boom of the later 1990s, and sketches a little bit of the history of the relevant data processing and data communications technologies. Sections III and IV present two of the simplest possible model that can handle the phenomena of a technological revolution like that we are going through, and shows that it predicts that it would be difficult not to have rapid productivity growth over the next decade. Section V briefly argues that there are important further considerations that tend to put more weight on the “optimistic” case. And section VI tries, in conclusion, to step back and provide a broader view.

II. The Pattern of Growth in the Later 1990s

Compare our use of information technology today with our predecessors' use of information technology half a century ago. 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

¹¹ See Crafts (2002).

702s, or DEC PDP-1s. The processing power of these machines averaged perhaps 10,000 machine instructions per second.

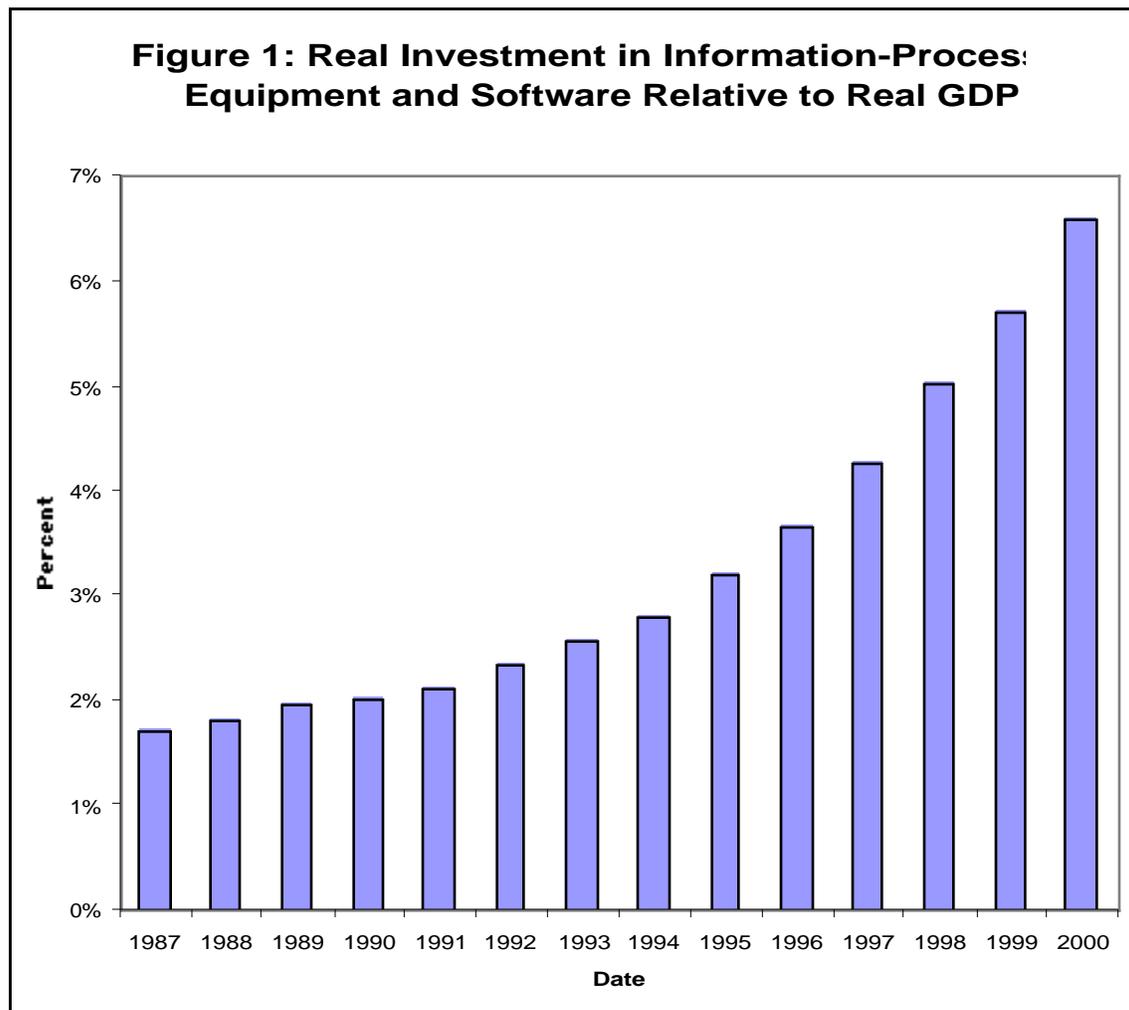
Today, talking rough orders of magnitude only, there are perhaps 300 million active computers in the world with processing power averaging several hundred million instructions per second. Two thousand computers times ten thousand instructions per second is twenty million. three hundred million computers times, say, three hundred million instructions/second is ninety quadrillion--a four-billion-fold increase in the world's raw automated computational power in forty years, an average annual rate of growth of 56 percent per year.

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 co-founder Gordon Moore noticed what has become Moore's Law--that improvements in semiconductor fabrication allow manufacturers to double the density of transistors on a chip every eighteen months. The scale of investment needed to make Moore's Law hold has grown exponentially along with the density of transistors and circuits, but Moore's Law has continued to hold, and engineers see no immediate barriers that will bring the process of improvement to a halt anytime soon.

Investment Spending

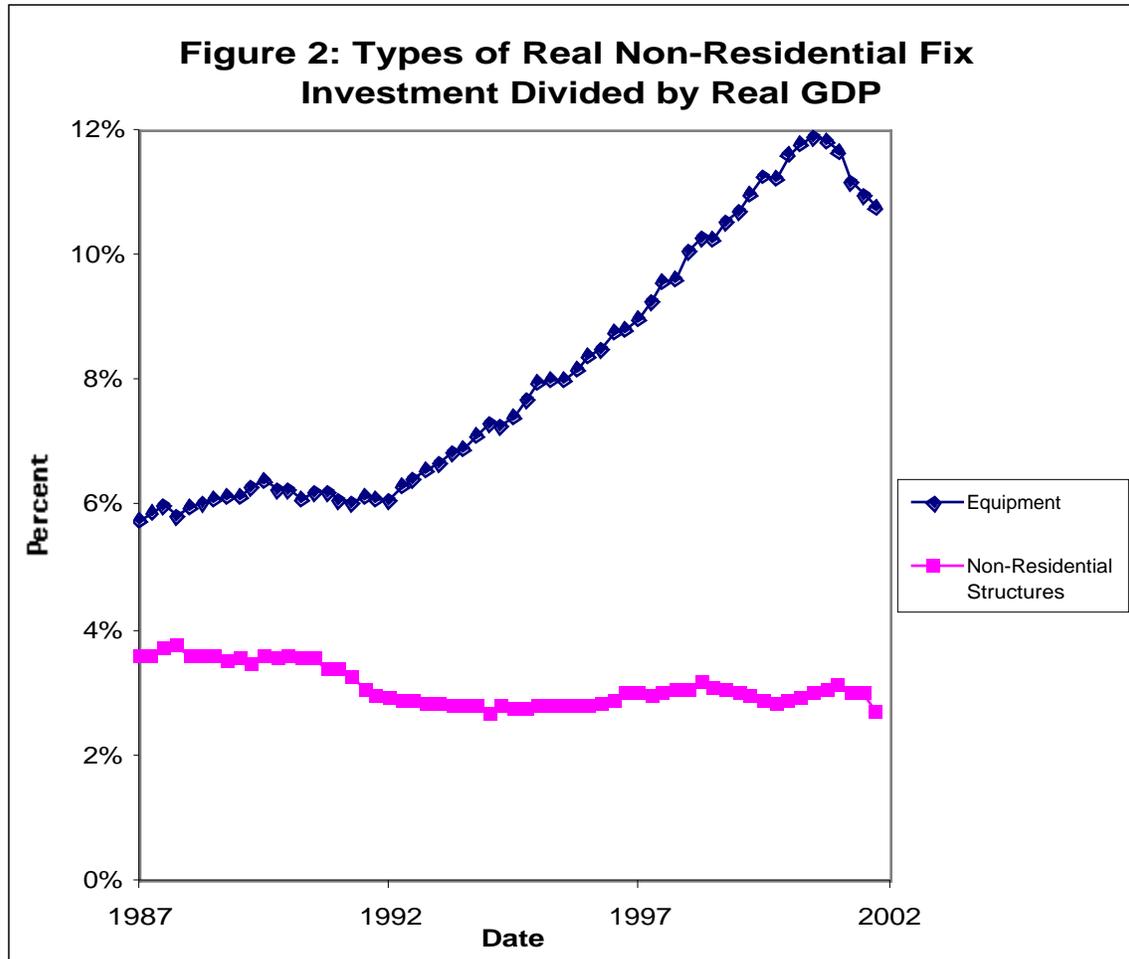
As the computer revolution proceeded, nominal spending on information technology capital rose from about one percent of GDP in 1960 to about two percent of GDP by 1980 to about three percent of GDP by 1990 to between five and six percent of GDP by 2000. All throughout this time, Moore's Law—the rule of thumb enunciated by Intel cofounder Gordon Moore that every twelve to eighteen months saw a doubling of the density of transistors that his and other companies could put onto a silicon wafer—meant that the real price of information technology capital was falling as well. As the nominal spending share of GDP spent on information technology capital grew at a rate of 5 percent per year, the price of data processing—and in recent decades data communications—equipment fell at a rate of between 10 and 15 percent per year as well. At chain-weighted real values constructed using 1996 as a base year, real investment in information technology equipment and software was an amount equal to 1.7 percent of real GDP in 1987. By 2000 it was an amount equal to 6.8 percent of real GDP.¹²

¹² Note that we can no longer think of these as “shares”: chain-weighted components of real GDP do not add up to the total of real GDP. For an excellent overview of what forms of addition and comparison are or are not legitimate using real chain-weighted values, see Karl Whelan (2000a).



Source: National Income and Product Accounts

The steep rise in real investment in information processing equipment (and software) drove a step 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.



Source: National Income and Product Accounts.

Macro Consequences

A naïve 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 percent per year, a 6 percentage-point rise in

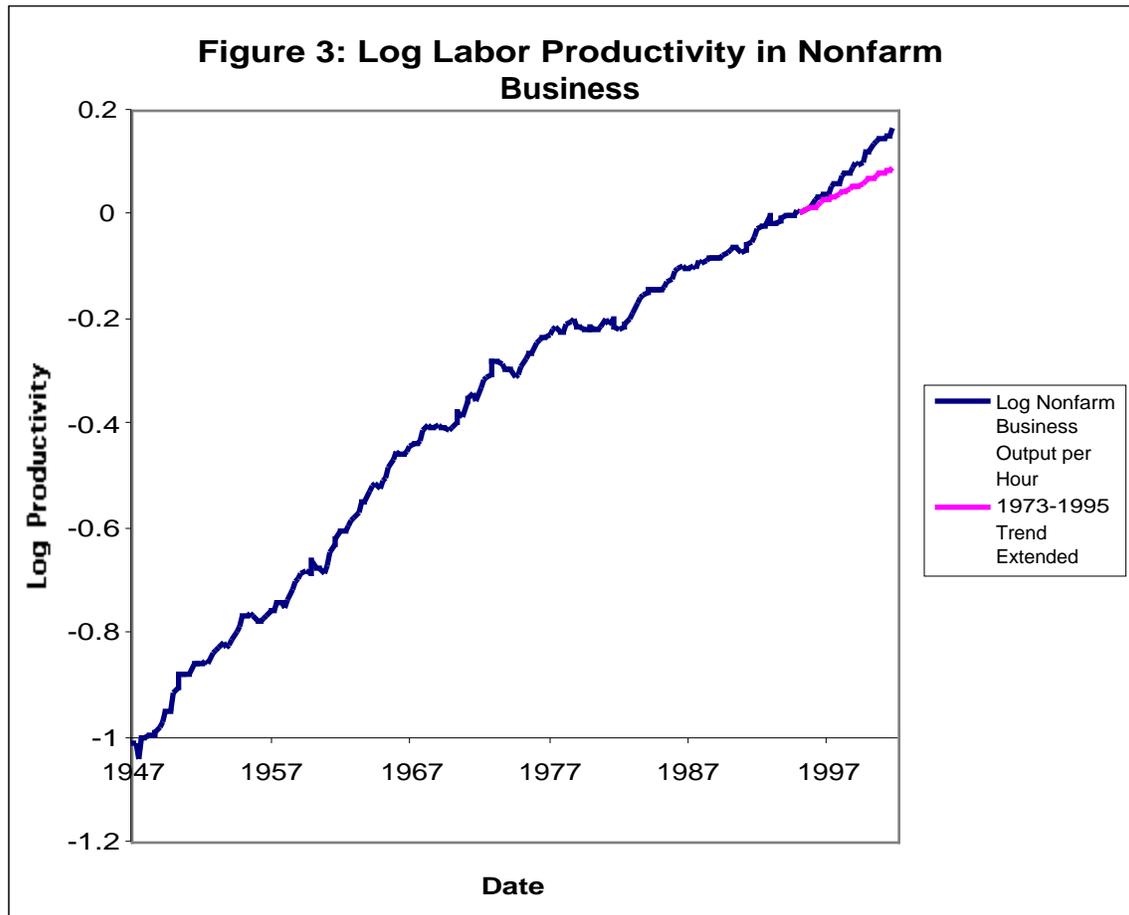
the investment share would be predicted to boost the rate of growth of real gross product, at least, by about 1 percentage point per year. And that is the same order of magnitude as the acceleration of economic growth seen in the second half of the 1990s.

Table 1: The Growth Acceleration of the Later 1990s

	1947:1- 1973:2	1973:2- 1995:1	1995:1- 2001:2	1995:1- 2001:4
Real Output per Hour Worked	2.88%	1.30%	2.80%	2.69%
Total Real GDP	3.91%	2.73%	4.21%	3.85%

Source: National Income and Product Accounts. Real GDP in chained 1996-base year dollars. Real Output per Hour Worked in the nonfarm business sector.

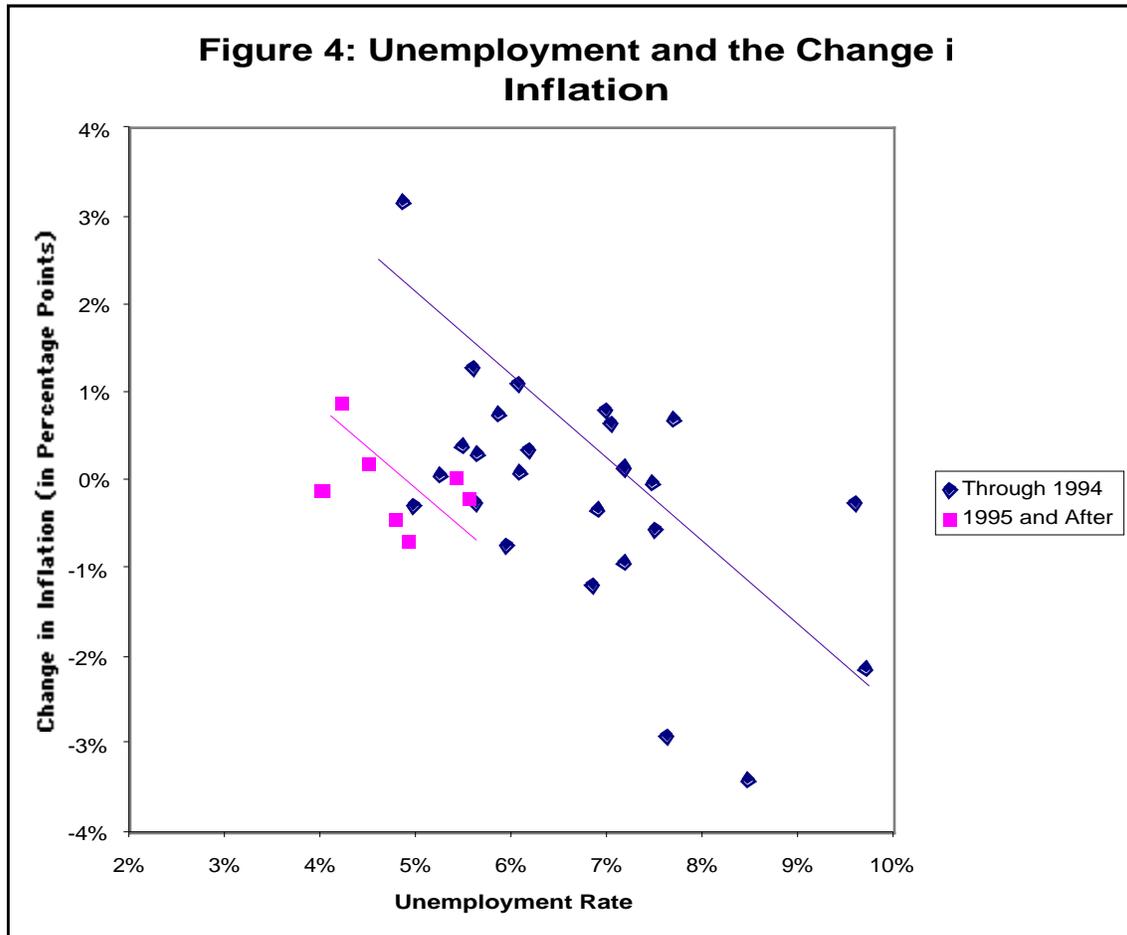
The acceleration in the growth rate of labor productivity and of real GDP in the second half of the 1990s effectively wiped out all the effects of the post-1973 productivity slowdown. The U.S. economy in the second half of the 1990s was, according to official statistics and measurements, performing as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



Source: National Income and Product Accounts

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 ten percent higher than one would have predicted back in 1995 by extrapolating the 1973 to 1995 trend. That such a large increase in the average level of productivity can be accumulated 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 to 1995 productivity slowdown.



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

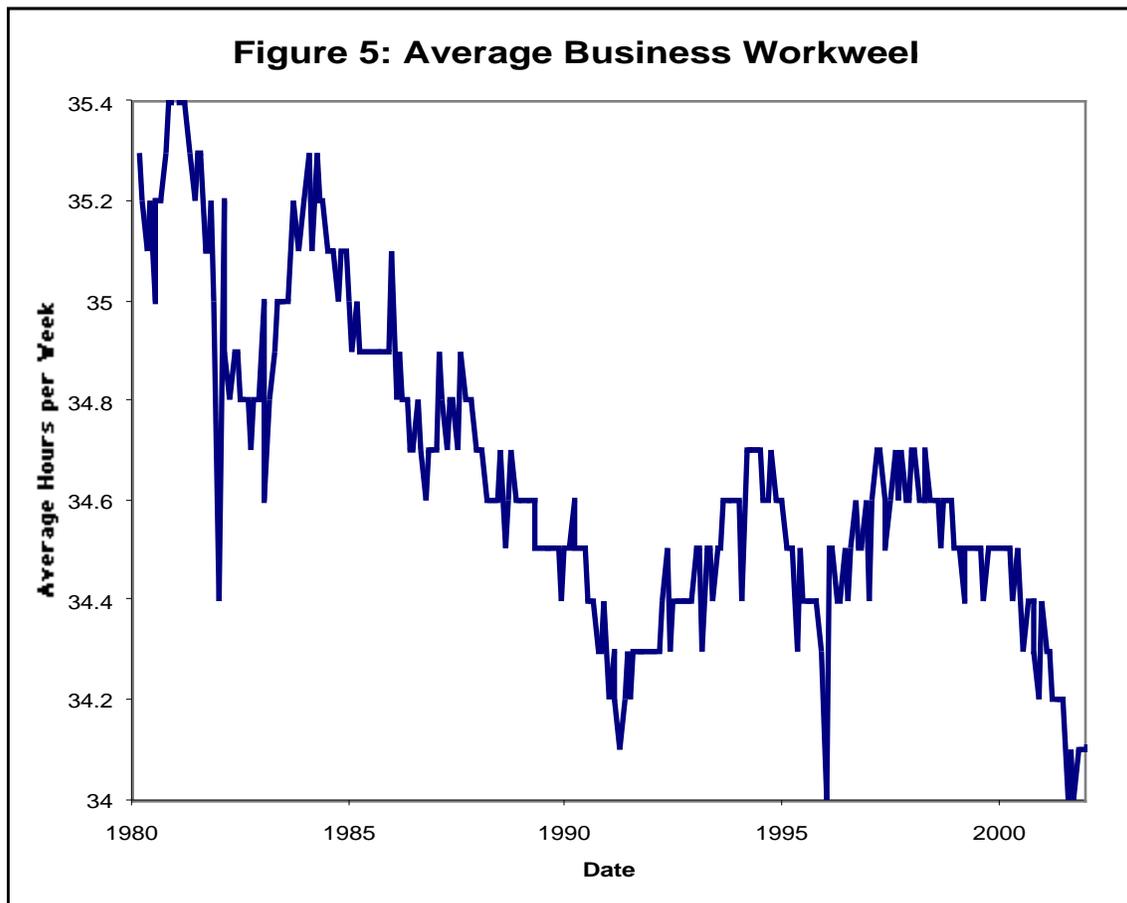
Cyclical Factors

Alongside the burst of growth in output per person-hour worked came significantly better labor market performance. The unemployment rate consistent with stable inflation, which had been somewhere between 6 and 7 percent of the labor force from the early 1980s into the early 1990s, suddenly fell to 5 percent or even lower in the late 1990s. All estimates of non-accelerating-inflation-rates-of-unemployment 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 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 two percentage-point decline in the unemployment rate would be associated with a 5 percent increase in output. Production functions would suggest that a two percentage-point decline in the unemployment rate would—after taking account of the effect of falling unemployment on the labor force and the differential impact of the

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

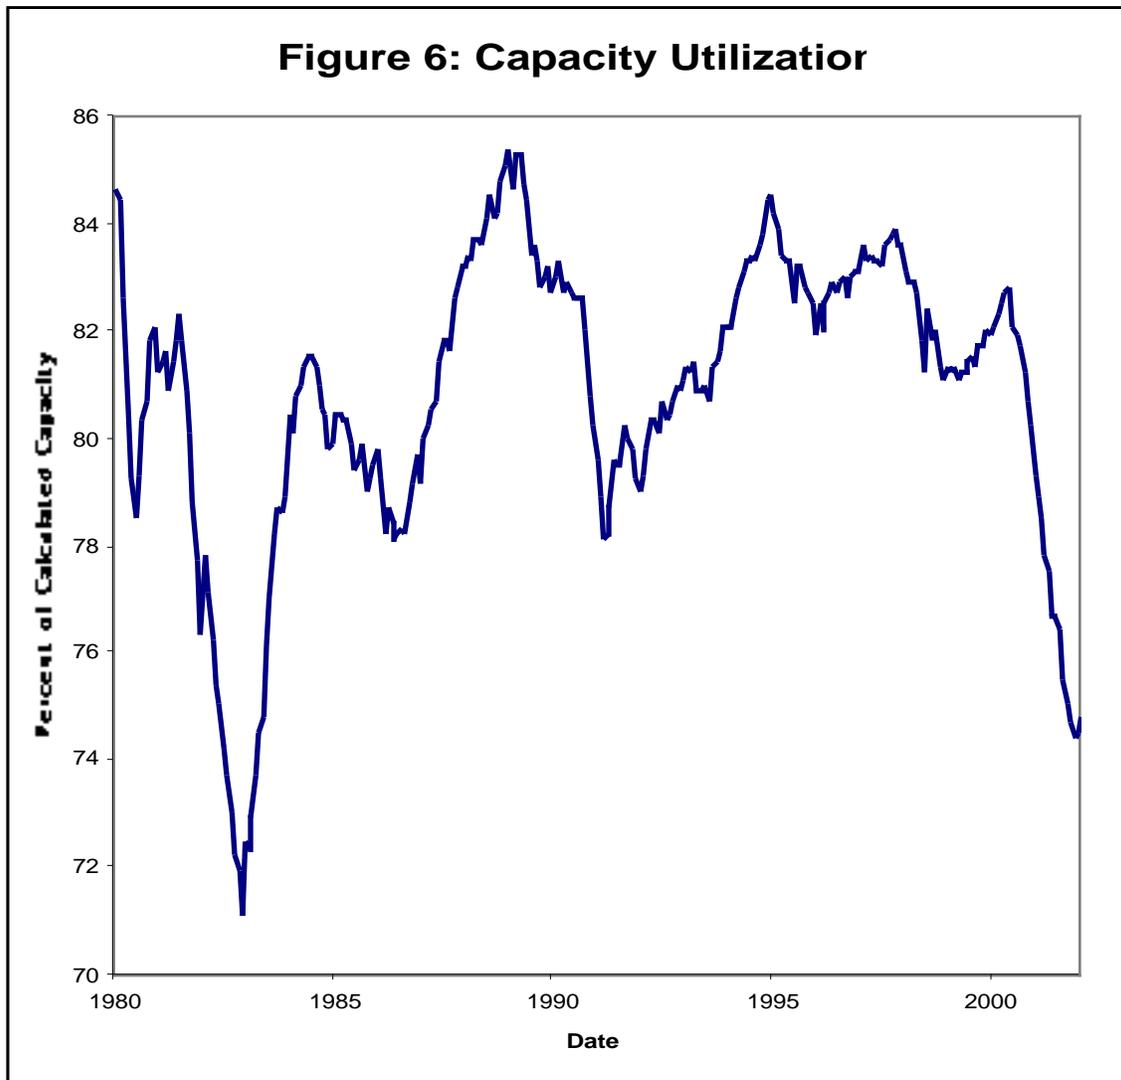
change in unemployment on the skilled and the educated—be associated with a roughly 1.5 percent increase in output.



Source: Bureau of Labor Statistics.

However, none of the other cyclical indicators suggested that the late-1990s economy was an unusually high-pressure economy. The average workweek was no higher in 2000

when the unemployment rate approached 4 percent than it had been in 1993 when the unemployment rate fluctuated between 6 and 7 percent.



Source: Federal Reserve.

Capacity utilization 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.

III. A Simple Model

Basic Theory

Suppose that the economy produces two types of output—regular goods, which we will denote by an r , and information technology 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 information technology capital and vice-versa. Thus:

$$(1) \quad Y_t = Y_t^r + (p_t^i \times Y_t^i)$$

Total output Y_t is equal to the output of regular goods Y_t^r plus $p_t^i \times Y_t^i$, the output of information technology capital multiplied by its current cost price.

Total output Y_t is itself determined by a standard production function:

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

$$(2) \quad Y_t = F(A_t, K_t^r, K_t^i, L_t)$$

Where A_t is the exogenous level of total factor productivity, K_t^r is the stock of “normal” capital, K_t^i is the stock of information technology capital, and L_t is the labor force.

Suppose further that because of ongoing technological revolutions the cost price of information technology is declining at a constant proportional rate of π .

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

$$(3) \quad \frac{d(\ln(Y^*_t))}{dt} = \left(\frac{\partial F}{\partial A}\right)\left(\frac{A}{Y}\right)\frac{d(\ln(A_t))}{dt} + \left(\frac{\partial F}{\partial K^r}\right)\left(\frac{K^r}{Y}\right)\frac{d(\ln(K^r_t))}{dt} + \left(\frac{\partial F}{\partial K^i}\right)\left(\frac{K^i}{Y}\right)\frac{d(\ln(K^i_t))}{dt} + \left(\frac{\partial F}{\partial L}\right)\left(\frac{L}{Y}\right)\frac{d(\ln(L_t))}{dt} + X_t^i\pi$$

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

Under the assumptions of constant returns to scale and competition, the $(\partial F/\partial K)(K/Y)$ 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 each of the

$(\partial F/\partial K)(K/Y)$ and like terms, and also normalize the total factor productivity term A, and thus rewrite (3) as:

$$(4) \quad \frac{d(\ln(Y^*_t))}{dt} = \frac{d(\ln(A_t))}{dt} + s_r \frac{d(\ln(K^r_t))}{dt} + s_i \frac{d(\ln(K^i_t))}{dt} + s_L \frac{d(\ln(L_t))}{dt} + X^i \pi$$

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

$$(5) \quad \frac{d(\ln(Y^*_t))}{dt} = a + s_L n + s_r \left(\frac{X^r Y}{K^r} - \delta^r \right) + s_i \left(\frac{X^i Y}{p^i_t K^i} - \delta^i \right) + X^i \pi$$

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

$$(6) \quad \frac{d(\ln(Y^*_t / L_t))}{dt} = a + s_r \left(\frac{X^r Y}{K^r} - (\delta^r + n) \right) + s_i \left(\frac{X^i Y}{p^i_t K^i} - (\delta^i + n) \right) + X^i \pi$$

Analysis

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

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

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:

$$(8) \quad X^i \pi$$

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 information technology capital more and more cheaply in terms of normal goods is a genuine improvement in productivity.

The first term:

$$(9) \quad s_i \left\{ \frac{\left(\frac{X^i}{p_t^i} \right) Y}{K^i} - (\delta^i + n) \right\}$$

is the contribution to the production of normal output from the net increase in information technology capital stock per worker. However the numerator is not the nominal share of GDP expended on information technology capital X^i , but the real share X^i/p_t^i .

And—because the cost price of information technology capital is falling at the rate π —a constant nominal expenditure share means that the real expenditure share relevant for the contribution of information technology 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 information technology 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 information technology 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

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)

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 nominal regular capital-output ratio κ_t^r and the nominal current-dollar value information technology capital-output ratio κ_t^i by:

$$(10) \quad \kappa_t^r = \frac{K_t^r}{Y_t}$$

$$(11) \quad \kappa_t^i = \frac{p_t^i K_t^i}{Y_t}$$

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 :

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

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

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

and then derive rates-of-change of these nominal capital-output ratios:

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

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

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

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

to obtain:

$$(17) \quad \frac{Y_t}{L_t} = A_t \left(\frac{1}{1-s_r-s_i} \right) (\kappa_t^r)^{\left(\frac{s_r}{1-s_r-s_i} \right)} (\kappa_t^i)^{\left(\frac{s_i}{1-s_r-s_i} \right)} (p_t^i)^{\left(\frac{-s_i}{1-s_r-s_i} \right)}$$

The dynamics of output per worker the economy can then be analyzed in terms of the (constant) proportional increase in total factor productivity A, the (constant) proportional decrease in the real cost price of information technology goods, and the dynamic evolution of the nominal capital-output ratios:

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

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

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

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

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

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

However, such a steady-state analysis is of dubious validity. The steady-state assumes constant nominal investment shares in information-technology 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 decade of the 1980s very far away from its pseudo-steady state: back then the GDP share of nominal spending on information technology investment was only forty percent of its current

value, and likewise for the share of national income attributable to the information-technology capital stock.

IV. A Demand-Side Model

A Model of Changing Demand Shares

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

$$(1) \quad \pi = \sigma(\pi_L) + (1-\sigma)(\pi_R)$$

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

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 total factor productivity 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.

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.

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_l equal to 1.0, set the heroic age of invention and innovation to a period 40 years long, and set the background level of total factor productivity growth π_R to 0.01 per year, one percent 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 one-tenth of a percent 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 the expenditure share has fallen below 1.0%. By the fortieth year the expenditure share 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 forty years total productivity in the economy as a whole is only 2.54% higher than had the leading sector 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' (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 grow rapidly from their original value of 2%. With a productivity growth wedge of 5% per year, the leading sector share of spending surpasses 10% by year 12, 30% by year 20, and reaches 89% by year 40. As the spending share of the leading sectors rise, 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: total factor productivity 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 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?

How Useful Will Computers Be?

What factors determine what the ultimate impact of these technologies will be? What is there that could interrupt a relatively 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 information technology capital. The second is a steep fall in the share of total nominal expenditure devoted to information

technology capital. And the third is a steep fall in the social marginal product of investment in information technology—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 the future is that none of these have happened: In 1991-1995 semiconductor production was half a percent of nonfarm business output; in 1996-2000 semiconductor production averaged 0.9 percent of nonfarm business output. Nominal spending on information technology capital rose from about one percent of GDP in 1960 to about two percent of GDP by 1980 to about three percent of GDP by 1990 to between five and six percent of GDP by 2000. Computer and semiconductor prices declined at 15-20 percent per year from 1991-1995 and at 25-35 percent per year from 1996-2000.

However, whether nominal expenditure shares will continue to rise in the end hinges on how useful data processing and data communications 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, not 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.

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 was 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 electro-mechanical tabulating machines. Businesses used computers to do the payroll, report-generating, and record-analyzing tasks that their own electro-mechanical 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 reservations 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 is instead as a what-if machine. The computer creates models of what-if: 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 America's 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.

V. Additional Considerations

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

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 total factor

productivity 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 total factor productivity 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 (2000) 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.

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 ten percent of national product in equipment this year and two percent the next is much worse than investing a steady six percent 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 naïve 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.

VI. Conclusion

The macroeconomics 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 information technology, 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 for one consumer and a different price for essentially the same good

for 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 economy 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 federal 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 economy 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 equal marginal cost. However, forcing prices to be equal to marginal cost cannot be sustained because then the fixed set-up costs are not covered. Relying on government subsidies to cover fixed set-up 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 administrative 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 butchershop. 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 of the future 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.

References

Robert Barro (1999), "Notes on Growth Accounting", *Journal of Economic Growth*, 4, 119- 137.

Martin Baily (2002), "The New Economy: Post-Mortem or Second Wind?" *Journal of Economic Perspectives* (forthcoming).

Susanto Basu, John Fernald, and Matthew Shapiro (2001), "Productivity Growth in the 1990s: Technology, Utilization, or Adjustment?" (Cambridge: NBER Working Paper 8359).

Timothy Bresnahan and Manuel Trajtenberg (1995), "General Purpose Technologies: 'Engines of Growth'?", *Journal of Econometrics* 65, pp. 83-108.

Erik Brynjolffson and Loren Hitt (2000), "Beyond Computation: Information Technology, Organizational Transformation and Business Performance", *Journal of Economic Perspectives*, 14(4), 23-48.

N.F.R. Crafts (2002), "The Solow Productivity Paradox in Historical Perspective" (London: LSE xerox).

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

Paul David (1991), "Computer and Dynamo: The Modern Productivity Paradox in a Not-Too-Distant Mirror", in OECD, *Technology and Productivity: The Challenge for Economic Policy*. (Paris: OECD).

Phyllis Deane and W.A. Cole (1962), *British Economic Growth, 1688-1959*. Cambridge: Cambridge University Press.

Walter Devine (1983), "From Shafts to Wires: Historical Perspective on Electrification", *Journal of Economic History*, 43, 347-372.

Stanley Fischer (1988), "Symposium on the Slowdown in Productivity Growth," *Journal of Economic Perspectives* 2:4 (Fall), pp. 3-7.

Robert Gordon (2002), "Technology and Economic Performance in the American Economy" (Cambridge: NBER Working Paper 8771).

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

Robert Gordon (2000b), "Interpreting the 'One Big Wave' in U.S. Long Term Productivity Growth" (Cambridge: NBER Working Paper 7752).

Jeremy Greenwood and Boyan Jovanovic (1999). "The Information Technology Revolution and the Stock Market," *American Economic Review* 89:2 (May), pp. 116-122.

Zvi Griliches (1988), "Productivity Puzzles and R & D: Another Nonexplanation," *Journal of Economic Perspectives* 2:4 (Fall), pp. 9-21.

C. Knick Harley (1999), "Reassessing the Industrial Revolution: A Macro View", in Joel Mokyr (ed.), *The British Industrial Revolution: An Economic Perspective*. Oxford: Westview Press, 160-205.

Dale Jorgenson (1988), "Productivity and Postwar U.S. Economic Growth," *Journal of Economic Perspectives* 2:4 (Fall), pp. 23-41

Dale Jorgenson and Kevin Stiroh (2000), "Raising the Speed Limit: US Economic Growth in the Information Age", *Brookings Papers on Economic Activity*, 1, 125-211.

Dale Jorgenson, Mun Ho, and Kevin Stiroh (2001), "Projecting Productivity Growth: Lessons from the U.S. Growth Resurgence" (Cambridge: Harvard University xerox).

Julie Kosterlitz (2002), "What Went Wrong?" *National Journal* (January 4, 2002).

- Paul Krugman (1994), *The Age of Diminished Expectations*, 4th ed. (Cambridge: MIT Press).
- David Landes (1969), *The Unbound Prometheus* (Cambridge: Cambridge University Press).
- McKinsey Global Institute (2001), "U.S. Productivity Growth, 1995-2000" (Washington: McKinsey).
- Steven Oliner and Daniel Sichel (2000), "The Resurgence of Growth in the Late 1990s: Is Information Technology the Story?", *Journal of Economic Perspectives*, 14(4), 3-22.
- Steven Oliner and Daniel Sichel (2002), "Information Technology and Productivity: Where Are We Now and Where Are We Going?" (Washington: Federal Reserve Board xerox).
- Nicholas Oulton (2001), "ICT and Productivity Growth in the United Kingdom" (London: Bank of England Working Paper 140).
- Paul Romer (1986), "Increasing Returns and Long Run Growth ", *Journal of Political Economy* 95, pp. 1002-1037.
- Paul Romer (1990), "Endogenous Technological Change", *Journal of Political Economy*, 98, S71-S102.
- Paul Schreyer (2000), "The Contribution of Information and Communication Technology to Output Growth: A Study of the G7 Countries", OECD STI Working Paper No 2000/2.
- Daniel Sichel (2001), "The Resurgence of US Growth in the Late 1990s: An Update and What Lies Ahead."
- Douglas Staiger, James H. Stock, and Mark W. Watson (1997), "The NAIRU, Unemployment and Monetary Policy," *Journal of Economic Perspectives* 11:1 (Winter), pp. 33-49.

Kevin Stiroh (1998), "Computers, Productivity, and Input Substitution", *Economic Inquiry*, 36, 175-191.

Kevin Stiroh (2001), "Information Technology and the US Productivity Revival: What Do the Industry Data Say ?" (New York: Federal Reserve Board of New York Working Paper No. 115).

Jack Triplett (1999), "The Solow Productivity Paradox: What Do Computers Do to Productivity?", *Canadian Journal of Economics*, 32, 309-334.

Nicholas von Tunzelmann (1978), *Steam Power and British Industrialization to 1860*. Oxford: Clarendon Press.

Karl Whelan (2001), "A Two Sector Approach to Modelling NIPA Data" (Washington: Federal Reserve).

Karl Whelan (2000a), "A Guide to the Use of Chain-Aggregated NIPA Data" (Washington: Federal Reserve xerox).

Karl Whelan (2000b), "Computers, Obsolescence, and Productivity", mimeo, Federal Reserve.