Explaining a Productive Decade

STEPHEN D. OLINER
Board of Governors of the Federal Reserve System

DANIEL E. SICHEL
Board of Governors of the Federal Reserve System

KEVIN J. STIROH
Federal Reserve Bank of New York

August 2007

We thank Martin Baily, John Fernald, Andrew Figura, Dale Jorgenson, Gregory Mankiw, participants at the March 2007 Brookings Panel conference, and participants at the March 2007 Productivity Meeting at the National Bureau of Economic Research for useful comments and discussions; we also thank John Roberts for providing code and assistance to implement his Kalman filter model, George Smith and Robert Yuskavage of the Bureau of Economic Analysis for help with the industry data, and David Byrne for assistance with semiconductor data. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of New York, the Board of Governors of the Federal Reserve System, or other staff members at either organization.

Abstract

This paper analyzes the sources of U.S. productivity growth in recent years using both aggregate and industry-level data. We confirm the central role for information technology (IT) in the productivity revival during 1995-2000 and show that IT played a significant, though smaller, role after 2000. Productivity growth after 2000 appears to have been boosted by industry restructuring and cost cutting in response to profit pressures, an unlikely source of future strength. In addition, the incorporation of intangible capital into the growth accounting framework takes some of the luster off the performance of labor productivity since 2000 and makes the gain during 1995-2000 look larger than in the official data. Finally, we examine the outlook for trend growth in labor productivity; our estimate, though subject to much uncertainty, is centered at 2-1/4 percent a year, faster than the lackluster pace that prevailed before 1995 but somewhat slower than the 1995-2006 average.

Productivity growth in the United States rose sharply in the mid-1990s, after a quarter century of sluggish gains. That pickup was widely documented, and a relatively broad consensus emerged that the speedup in the second half of the 1990s was importantly driven by information technology (IT).[1] After 2000, however, the economic picture changed dramatically, with a sharp pullback in IT investment, the collapse in the technology sector, the terrorist attacks of September 11, 2001, and the 2001 recession. Given the general belief that IT was a key factor in the growth resurgence in the mid-1990s, many analysts expected that labor productivity growth would slow as IT investment retreated after 2000. Instead labor productivity accelerated further over the next several years. More recently, however, the pace of labor productivity growth has slowed considerably.

In light of these developments, researchers and other commentators have been intensely interested in the course of productivity growth since 2000. Distinguishing among the possible explanations for the continued strength in productivity growth is challenging, because much of that strength appeared in measured multifactor productivity (MFP), the unexplained residual in the standard growth accounting setup. Nevertheless, potential explanations can be divided into two broad categories: those centered on IT and those unrelated or only loosely related to IT.

The simplest IT-centered story—that rapid technological progress in the production of IT and the induced accumulation of IT capital raised productivity growth—does not work for the period after 2000, because the contributions to growth from both the production and the use of IT declined. A second IT-related story that has received a great deal of attention is that IT

---

1. See *Economic Report of the President 2001*, Basu, Fernald, and Shapiro (2001), Brynjolffson and Hitt (2000, 2003), Jorgenson and Stiroh (2000), Jorgenson, Ho, and Stiroh (2002, 2005), and Oliner and Sichel (2000, 2002). In these papers IT refers to computer hardware, software, and communications equipment. This category often also is referred to as information and communications technology, or ICT. For industry-level evidence supporting the role of IT in the productivity resurgence, see Stiroh (2002b). For an interpretation of the industry evidence that puts less emphasis on IT, see Bosworth and Triplett (2007) and McKinsey Global Institute (2001).

investment proxies for complementary investments in intangible capital, and a growing body of research has highlighted the important role played by such intangibles.[2] A third IT-related story identifies IT as a general-purpose technology that spurs further innovation over time in a wide range of industries, ultimately boosting growth in MFP.[3] Because this process takes time, the gains in MFP observed since 2000 could reflect the follow-on innovations from the heavy investment in IT in the second half of the 1990s.

Another broad set of explanations highlights forces not specific to IT. Gains in labor productivity since 2000 could have been driven by fundamental technological progress outside of IT production, as implied by the strong growth in MFP in other sectors.[4] Alternatively, the robust advance in labor productivity could reflect broader macroeconomic factors such as normal cyclical dynamics, a decline in adjustment costs after 2000 as investment spending dropped back, greater-than-usual business caution in hiring and investment, or increased competitive pressures on firms to restructure, cut costs, raise profits, and boost productivity. The profit-driven cost-cutting hypothesis, in particular, has received considerable attention in the business press.[5]

In this paper we try to sort out these issues using both aggregate and industry-level data.[6] We investigate four specific questions. First, given the latest data and some important extensions

2. See Corrado, Hulten, and Sichel (2005, 2006), Brynjolfsson and Hitt (2003), Bresnahan, Brynjolfsson, and Hitt (2002), Basu and others (2004), Black and Lynch (2001, 2004), and Nakamura (1999, 2001, 2003). The National Income and Product Accounts (NIPAs) exclude virtually all intangibles other than software, although the Bureau of Economic Analysis, which produces the NIPA data, recently released a satellite account for scientific research and development; see Okubo and others (2006).
3. Bresnahan and Trajtenberg (1995) were the first to write about IT as a general-purpose technology. Also, see Organization for Economic Cooperation and Development (2000), Schreyer (2000), van Ark (2000), Basu and others (2004), and Basu and Fernald (2007).
4. Jorgenson, Ho, and Stiroh (2007); Bosworth and Triplett (2007).
5. See Gordon (2003), Baily (2003), Schweitzer (2004), and Stiroh (2006a). For references to the business press, see Gordon (2003) and Stiroh (2006a).
6. Several other researchers have examined industry data, including Baily and Lawrence (2001), Stiroh (2002b), Nordhaus (2002b), Corrado and others (2007), and Bosworth and Triplett (2007). For references to the literature on industry-level data in Europe, see van Ark and Inklaar (2005).

to the standard growth accounting framework, is an IT-centered story still the right explanation for the resurgence in productivity growth over 1995-2000, and does IT play a significant role when considering the entire decade since 1995? Second, what accounts for the continued strength in productivity growth after 2000? Third, how has investment in intangible capital influenced productivity developments? Finally, what are the prospects for labor productivity growth in coming years?

Our analysis relies in part on neoclassical growth accounting, a methodology that researchers and policymakers have used for many years to gain insights into the sources of economic growth. Notably, the Council of Economic Advisers, the Congressional Budget Office, and the Federal Reserve Board routinely use growth accounting as part of their analytical apparatus to assess growth trends.[7]

Of course, growth accounting is subject to limitations, and in recent years many analysts have leveled critiques at this methodology. For example, the standard neoclassical framework does not explicitly account for adjustment costs, variable factor utilization, deviations from perfect competition and constant returns to scale, outsourcing and offshoring, management expertise, or the intangibles that are omitted from published data. In addition, researchers have raised a host of measurement issues that could affect the standard framework.[8] It is well beyond the scope of this paper to deal with all of these critiques, but we augment the standard framework

7. See *Economic Report of the President 2007,* Congressional Budget Office (2007a, 2007b), and the latest available transcripts of the meetings of the Federal Open Market Committee (FOMC; Board of Governors of the Federal Reserve System, 2001). The 2001 FOMC transcripts show that staff presentations on the economic outlook featured growth accounting in the discussion of productivity trends. Private sector analysts also rely on growth accounting; see, for example, Global Insight (2007) and Macroeconomic Advisers (2007).
8. Much has been written about the link between management expertise and productivity, including Bloom and van Reenen (2006), McKinsey Global Institute (2001), and Farrell, Baily, and Remes (2005). Gordon (2003) and Sichel (2003) provide reasons why offshoring and hours mismeasurement may have had a relatively limited effect on labor productivity growth, whereas Houseman (2007) argues that these factors could have had a significant effect in the U.S. manufacturing sector. For a discussion of measurement issues related to the pace of technical progress in the semiconductor industry, see Aizcorbe, Oliner, and Sichel (2006). For further discussion of issues related to critiques of the neoclassical framework, see Congressional Budget Office (2007b).

to account for some of the most salient ones. In particular, we take on board time-varying utilization of inputs, adjustment costs for capital, and intangibles. Our intention is to broaden the standard framework to get a fuller view of productivity developments during the past decade.

Briefly, our answers to the four questions we pose are as follows. Both the aggregate and the industry-level results indicate that IT was indeed a key driver of the pickup in labor productivity growth over 1995-2000. IT also is a substantial contributor to labor productivity growth over the full decade since 1995, although its contribution is smaller after 2000. In the aggregate data, this conclusion stands even after accounting for variable factor utilization, adjustment costs, and intangible capital.

Regarding the continued strength in labor productivity growth since 2000 in the published data, our answer has a number of elements. As a matter of growth accounting arithmetic, the smaller—although still sizable—contribution of IT after 2000 was more than offset by several factors, the most important being faster MFP growth outside the IT-producing sector. Just as the aggregate data highlight different sources of productivity growth during 1995-2000 than since 2000, so do the industry data. The industry composition of labor productivity growth across these periods shifted significantly, and we report evidence that IT capital was linked to changes in industry productivity growth in the 1990s but not in the period since 2000.

The industry data also suggest that the rapid post-2000 productivity gains were due, at least in part, to restructuring and cost cutting in some industries as highlighted by Robert Gordon.[9] In particular, those industries that saw the sharpest declines in profits from the late 1990s through 2001 also tended to post the largest gains in labor productivity in the early 2000s. Because these restructuring-induced advances probably were one-time events (and could be reversed), they are unlikely to be a source of ongoing support to productivity growth.

9. Gordon (2003).

In addition, the industry evidence indicates that reallocations of both material and labor inputs have been important contributors to labor productivity growth since 2000, a point that Barry Bosworth and Jack Triplett also note.[10] Although it is difficult to pin a precise interpretation on the reallocation results, the importance of these reallocations could be viewed as evidence that the flexibility of the U.S. economy has supported aggregate productivity growth in recent years by facilitating the shifting of resources among industries.

The incorporation of intangibles into the aggregate growth accounting framework takes some of the luster off the performance of labor productivity since 2000 and makes the gains in the 1995-2000 period look better than in the published data. In addition, the step-up after 2000 in MFP growth outside the IT-producing sector is smaller after accounting for intangibles than in the published data. Thus any stories tied to a pickup in MFP growth (such as IT as a general-purpose technology) may apply to the entire decade since 1995 and not simply to recent years. This framework also implies that intangible investment has been quite sluggish since 2000, coinciding with the soft path for IT capital spending. All else equal, this pattern could be a negative for labor productivity growth in the future to the extent that these investments are seed corn for future productivity gains.

Finally, our analysis of the prospects for labor productivity growth highlights the wide range of possible outcomes. We report updated estimates of trend growth in labor productivity from a Kalman filter model developed by John Roberts;[11] these results generate a 2-standard-error confidence band extending from 1¼ percent to 3¼ percent at an annual rate, with a point estimate of 2¼ percent. In addition, we solve for the steady-state growth of labor productivity in a multisector model under a range of conditioning assumptions. This machinery also suggests a

10. Bosworth and Triplett (2007).
11. Roberts (2001).

wide range of outcomes, extending from about 1½ percent to just above 3 percent, with a midpoint of 2¼ percent. Notwithstanding the wide band of uncertainty, these estimates are consistent with productivity growth remaining significantly above the pace that prevailed in the twenty-five years before 1995, but falling short of the very rapid gains recorded over the past decade.

The paper is organized as follows. The next section reviews the aggregate growth accounting framework and presents baseline results that account for variable factor utilization and adjustment costs. The section that follows uses the approach of Susanto Basu and coauthors to generate time series for intangible investment and capital services and presents growth accounting results for the augmented framework.[12] This approach complements that in the 2005 and 2006 papers by Carol Corrado, Charles Hulten, and Daniel Sichel, who also developed time series of intangible investment and capital and incorporated those estimates into a standard growth accounting framework. We then turn to the industry data to supplement the insights that can be drawn from the aggregate data. Finally, we discuss the outlook for productivity growth and present some brief conclusions.

Aggregate Growth Accounting: Analytical Framework and Baseline Results

We use an extension of the growth accounting framework developed by Oliner and Sichel to analyze the sources of aggregate productivity growth in the United States.[13] That framework was designed to measure the growth contributions from the production and use of IT capital, key factors that emerged in the second half of the 1990s. The framework has some limitations, however. It excludes intangible capital, which has received much attention in recent research on the sources of productivity gains. It also imposes the strict neoclassical assumption

12. Basu and others (2004).
13. Oliner and Sichel (2000, 2002).

of a frictionless economy and thus abstracts from cyclical influences on productivity growth and from the effects of adjustment costs arising from the installation of new capital goods.

The growth accounting framework in this paper incorporates all of these considerations. We meld the original Oliner-Sichel model with the treatment of adjustment costs and cyclical factor utilization developed by Basu, John Fernald, and Matthew Shapiro.[14] In addition, we take account of intangible capital by drawing on the model of Basu, Fernald, Nicholas Oulton, and Sylaja Srinivasan.[15]

Analytical Framework

The model that underlies our analytical framework includes six sectors. Four of these produce the final nonfarm business output included in the National Income and Product Accounts (NIPAs): computer hardware, software, communications equipment, and a large non-IT-producing sector. The NIPAs omit production of virtually all intangible capital other than software. Our model accounts for this capital by adding a fifth final-output sector that produces the intangible assets excluded from the NIPAs. In addition to the five final-output sectors, our model includes a sector that produces semiconductors, which are either consumed as an intermediate input by the final-output sectors or exported to foreign firms. To focus on the role of semiconductors in the economy, the model abstracts from all other intermediate inputs.

Following BFS, we allow the length of the workweek, labor effort, and the utilization of capital to vary over time. We also assume that the installation of new capital diverts resources from the production of market output. As in BFS, these adjustment costs depend on the amount of investment relative to existing capital. Boosting the ratio of investment to capital increases the

---

14. Basu, Fernald, and Shapiro (2001; hereafter BFS)
15. Basu and others (2004; hereafter BFOS).

fraction of output that is lost to adjustment costs.[16] To complete the model specification, we assume that the production function in every sector exhibits constant returns to scale and that the economy is perfectly competitive.[17]

Given this model, the appendix shows that growth in aggregate labor productivity can be expressed as

$$(1) \qquad A\dot{L}P \equiv \dot{V} - \dot{H} = \sum_j \left(\alpha_j^K - \phi_j\right)\left(\dot{K}_j - \dot{H}\right) + \alpha^L \dot{q} + M\dot{F}P,$$

where a dot over a variable signifies the growth rate of that variable, V is aggregate value added in nonfarm business, H is aggregate hours worked, K_j is the aggregate amount of type-j capital used in the nonfarm business sector, α^L and α_j^K are, respectively, the income shares for labor and each type of capital, ϕ_j is the adjustment cost elasticity of output with respect to type-j capital, q is an index of labor quality, and *MFP* denotes multifactor productivity. The various types of capital include computer hardware, software, communications equipment, other tangible capital, and intangible capital other than software; each type of capital is produced by the corresponding final-output sector in our model. Except for the adjustment cost effect captured by ϕ_j, equation 1 is a standard growth decomposition. It expresses growth in labor productivity as the sum of the contribution from the increase in capital per hour worked (capital deepening), the contribution from the improvement in labor quality, and growth in aggregate MFP.[18]

---

16. Although BFS also include adjustment costs for labor in their model, they zero out these costs in their empirical work. We simply omit labor adjustment costs from the start. For additional discussion of capital adjustment costs and productivity growth, see Kiley (2001).

17. The results in BFS and in Basu, Fernald, and Kimball (2006) strongly support the assumption of constant returns for the economy as a whole. We invoke perfect competition as a convenience in a model that already has many moving parts.

18. The weight on $\dot{K}_j - \dot{H}$ represents the output elasticity of type-j capital. In the case without adjustment costs, $\phi_j = 0$, and so the income share α_j^K proxies for this output elasticity. However, in the presence of adjustment costs, the first-order condition for the optimal choice of capital yields the more general result shown in equation 1. In effect, the income share captures both the direct contribution of capital to production and the benefit of having an

Aggregate MFP growth, in turn, equals a share-weighted sum of the sectoral MFP growth rates:

$$(2) \qquad M\dot{F}P = \sum_i \mu_i M\dot{F}P_i + \mu_S M\dot{F}P_S,$$

where S denotes the semiconductor sector and i indexes the final-output sectors in our model (listed above). The weight for each sector equals its gross output divided by aggregate value added. These are the usual Domar weights that take account of the input-output relationships among industries.[19] Equation 2 has the same structure as its counterpart in an earlier paper by Oliner and Sichel.[20] The only formal difference is that including intangible capital increases the number of final-output sectors from four to five.[21]

Finally, the sectoral MFP growth rates in equation 2 can be expressed as

$$(3) \qquad M\dot{F}P_i = \xi_i \dot{W}_i - \sum_j \phi_{j,i}\left(\dot{I}_{j,i} - \dot{K}_{j,i}\right) + \dot{z}_i$$

for the final-output sectors and

$$(4) \qquad M\dot{F}P_S = \xi_S \dot{W}_S - \sum_j \phi_{j,S}\left(\dot{I}_{j,S} - \dot{K}_{j,S}\right) + \dot{z}_S$$

for semiconductor producers, where the ξ's represent the elasticity of sectoral output with respect to the workweek (W), the I's and K's denote sectoral investment and capital services for each type of capital, the ϕ's represent the sectoral adjustment cost elasticities for each type of capital,

---

extra unit of capital to absorb adjustment costs. The weight in equation 1 nets out the portion of the income share that relates to adjustment costs, as this effect is embedded in the MFP term discussed below.

19. Domar (1961).

20. Oliner and Sichel (2002).

21. In contrast to the expression for aggregate MFP growth in BFS, equation 2 contains no terms to account for reallocations of output, labor, or capital across sectors. The particularly clean form of equation 2 arises, in large part, from our assumption of constant returns to scale and the absence of adjustment costs for labor (which implies that competitive forces equate the marginal product of labor in all sectors). In addition, we have assumed that any wedge between the shadow value of capital and its user cost owing to adjustment costs is the same in all sectors. Given this assumption, reallocations of capital across sectors do not affect aggregate output.

and the z's represent the true level of technology. All of the ξ's and ϕ's take positive values.

In the BFS model that we adopt, firms vary the intensity of their factor use along all margins simultaneously, which makes the workweek a sufficient statistic for factor utilization in general. Lengthening the workweek boosts measured MFP growth in equations 3 and 4 as firms obtain more output from their capital and labor. Regarding adjustment costs, faster growth of investment spending relative to that of capital depresses measured MFP growth as firms divert resources from producing market output to installing new plant and equipment. The effects of factor utilization and adjustment costs drive a wedge between measured MFP growth and the true pace of improvement in technology $\dot{z}$.

Data, Calibration, and Measurement Issues

This section provides a brief overview of the data used for our aggregate growth accounting, discusses the calibration of key parameters, and addresses some important measurement issues.[22] The national accounts data that we discuss here exclude virtually all forms of intangible capital except for investment in computer software. We defer the consideration of intangible capital until the next section.

Our dataset represents an up-to-date reading on productivity developments through 2006 based on data available as of the end of March 2007. We rely heavily on the dataset assembled by the Bureau of Labor Statistics (BLS) for its estimates of MFP in the private nonfarm business sector. This dataset extended through 2005 at the time we conducted the analysis for this paper. We extrapolated the series required for our framework through 2006, drawing largely on corresponding series in the NIPAs.

22. For details on data sources, see the data appendix to Oliner and Sichel (2002).

To calculate the income share for each type of capital in our framework, we follow the BLS procedure that distributes total capital income across assets by assuming that each asset earns the same rate of return net of depreciation.[23] This is the same method used by Oliner and Sichel and by Jorgenson, Ho, and Stiroh.[24] Consistent with the standard practice in the productivity literature, we allow these income shares to vary year by year.[25]

These data and procedures generate a series for aggregate MFP growth via equation 1. Given this series as a top-line control, we estimate MFP growth in each sector with the "dual" method employed by various researchers in the past.[26] This method uses data on the prices of output and inputs, rather than their quantities, to calculate sectoral MFP growth. We opt for the dual approach because the sectoral data on prices are available on a more timely basis than the corresponding quantity data. Roughly speaking, the dual method compares the rate of change in a sector's output price with that of its input costs. Sectors in which prices fall quickly compared with their input costs are estimated to have experienced relatively rapid MFP growth.[27]

The expression that links aggregate and sectoral MFP growth (equation 2) involves the Domar weight for each sector, the ratio of the sector's gross output to aggregate value added. For

---

23. The weight on the capital deepening term in equation 1 for type-j capital equals its income share minus its adjustment cost elasticity. As discussed below, empirical estimates of these asset-specific elasticities are not available, which forces us to approximate the theoretically correct weights. Note that the weights on the capital deepening terms in equation 1 sum to one minus the labor share under constant returns to scale. We replace the theoretically correct weights with standard income-share weights that also sum to one minus the labor share. This approximation attaches the correct weight to aggregate capital deepening but may result in some misallocation of the weights across asset types.
24. Oliner and Sichel (2000, 2002); Jorgenson, Ho, and Stiroh (2002, 2007).
25. Year-by-year share weighting embeds the implicit assumption that firms satisfy the static first-order condition that equates the marginal product of capital with its user cost. Strictly speaking, this assumption is not valid in the presence of adjustment costs, as noted by BFS and by Groth, Nunez, and Srinivasan (2006). Both of those studies replace the year-by-year share weights with the average shares over periods of five years or more, in an effort to approximate a steady-state relationship that might be expected to hold on average over longer periods. We found, however, that our results were little changed by replacing year-by-year shares with period-average shares. Accordingly, we adhere to the usual share weighting practice in the literature.
26. Jorgenson and Stiroh (2000), Jorgenson, Ho, and Stiroh (2002, 2007), Oliner and Sichel (2000, 2002), and Triplett (1996), among others.
27. Oliner and Sichel (2002) give a nontechnical description of the way in which we implement the dual method, and the appendix to this paper provides the algebraic details.

the four NIPA-based final-output sectors, gross output simply equals the value of the sector's final sales, which we estimate using data from the Bureau of Economic Analysis (BEA). For the semiconductor sector we calculate gross output based on data from the Semiconductor Industry Association as well as data constructed by Federal Reserve Board staff to support the Federal Reserve's published data on U.S. industrial production.

The final step is to calculate the influence of adjustment costs and factor utilization on the growth of both aggregate and sectoral MFP. In principle, we could use equations 3 and 4 to calculate the effects at the sectoral level and then aggregate those effects using equation 2. However, as equations 3 and 4 show, this bottom-up approach requires highly disaggregated data on investment and the workweek and equally disaggregated output elasticities with respect to adjustment costs and the workweek (the ϕ's and the ξ's). Unfortunately, estimates of the required sectoral elasticities are not available.

To make use of readily available estimates, we work instead from the top down. That is, we model the effects of adjustment costs and the workweek for the nonfarm business sector as a whole and then distribute the aggregate effects across sectors. Let $\dot{W}$ and ξ denote, respectively, the percentage change in the workweek for aggregate nonfarm business and the elasticity of nonfarm business output with respect to this aggregate workweek. Then the workweek effect for aggregate nonfarm business equals $\xi\dot{W}$. Similarly, we measure the aggregate effect of adjustment costs as $\phi(\dot{I}-\dot{K})$, where $\dot{I}, \dot{K}$, and ϕ denote, respectively, growth in aggregate real investment spending, growth in aggregate real capital services, and the aggregate adjustment cost elasticity. To complete the top-down approach, we assume that the adjustment cost and workweek effects are uniform across sectors. Under this assumption, the top-down version of equations 2 through 4 is as follows (starting with the sectoral equations):

(5) $$M\dot{F}P_i = \frac{1}{\bar{\mu}}\left(\xi\dot{W} - \phi(\dot{I} - \dot{K})\right) + \dot{z}_i$$

(6) $$M\dot{F}P_S = \frac{1}{\bar{\mu}}\left(\xi\dot{W} - \phi(\dot{I} - \dot{K})\right) + \dot{z}_S$$

(7) $$M\dot{F}P = \sum_i \mu_i M\dot{F}P_i + \mu_S M\dot{F}P_S = \xi\dot{W} - \phi(\dot{I} - \dot{K}) + \sum_i \mu_i \dot{z}_i + \mu_S \dot{z}_S,$$

where $\bar{\mu} \equiv \sum_i \mu_i + \mu_S$. One can easily verify that the second equality holds in equation 7 by substituting for MFP_i and MFP_S from equations 5 and 6. Equations 5 through 7 serve as our empirical counterpart to equations 2 through 4.

We follow BFS in specifying ξ, W, and ϕ. Starting with the workweek effect, we specify the aggregate elasticity ξ to be a weighted average of BFS's sectoral estimates of ξ for durable manufacturing, nondurable manufacturing, and nonmanufacturing. Using weights that reflect current-dollar output shares in these sectors, we obtain an aggregate value of ξ equal to 1.24. To measure the workweek itself, we use the BLS series for production or nonsupervisory workers from the monthly survey of establishments. Because the workweek in equations 5 through 7 is intended to measure cyclical variation in factor use, we detrend the log of this monthly series with the Hodrick-Prescott filter (with λ = 10,000,000 as in BFS) and use the detrended series to calculate $\dot{W}$ on an annual basis.

With regard to adjustment costs, we set the output elasticity ϕ equal to 0.035.[28] This elasticity is based on estimates of capital adjustment costs by Shapiro.[29] More recent studies provide estimates of adjustment costs on both sides of ϕ = 0.035. Robert Hall estimates capital adjustment costs in an Euler equation framework similar to Shapiro's but uses more-

28. BFS used a larger value for ϕ, 0.05, but subsequently corrected some errors that had affected that figure. These corrections caused the value of ϕ to be revised to 0.035.
29. Shapiro (1986).

disaggregated data and a different set of instruments for estimation.[30] Hall cannot reject the hypothesis that $\phi = 0$. In contrast, Charlotta Groth, using industry-level data for the United Kingdom, estimates ϕ to be about 0.055.[31] The divergent results in these studies highlight the uncertainty surrounding estimates of capital adjustment costs but do not suggest the need to move away from a baseline estimate of $\phi = 0.035$. We apply this elasticity to the difference between the growth rates of aggregate real business fixed investment from the NIPAs and the corresponding capital services series $(\dot{I} - \dot{K})$.

To summarize, we use annual data from BEA and BLS through 2006 to implement the aggregate growth accounting framework in equation 1. This framework yields an annual time series for aggregate MFP growth. We then use the dual method to allocate this aggregate MFP growth across sectors. Finally, we calculate the effects of adjustment costs and changes in factor utilization on both aggregate and sectoral MFP growth, drawing heavily on parameter values reported by BFS.

Results

Table 1 presents our decomposition of labor productivity growth in the nonfarm business sector using the published data described above. These data exclude intangible capital other than business investment in software, which, again, is already treated as an investment good in the NIPAs. The next section fully incorporates intangible capital into our measurement system and presents an augmented set of growth accounting results.

Focusing first on the published data, table 1 shows that average annual growth in labor productivity picked up from about 1.5 percent a year during 1973-95 to about 2.5 percent during the second half of the 1990s and then rose further, to more than 2.8 percent, in the period after

30. Hall (2004).
31. Groth (2005).

2000. Our results indicate that an important part of the initial acceleration (about 0.6 percentage point of the total speedup of just over 1 percentage point) reflected the greater use of IT capital. In addition, growth of MFP rose notably in the IT-producing sectors, with an especially large increase for producers of semiconductors. The pickup for the semiconductor sector mirrors the unusually rapid decline in semiconductor prices from 1995 to 2000, which the model interprets as a speedup in MFP growth.[32] The last line of the table shows that, all told, IT capital deepening and faster MFP growth for IT producers more than accounted for the total speedup in labor productivity growth during 1995-2000. These results confirm that the IT-centric story for the late 1990s holds up after incorporating the latest vintage of data and extending the framework to account for adjustment costs and utilization effects.

The table also quantifies the influence of adjustment costs and changes in utilization during this period (the two lines under "Growth of MFP"). These two factors, on net, do not explain any of the upward swing in MFP growth from 1973-95 to 1995-2000, consistent with the results in BFS. Although the greater utilization of capital and labor had a positive effect on MFP growth during 1995-2000, this influence was offset by the negative effect from the higher adjustment costs induced by the investment boom of that period.

Table 1 tells a sharply different story for the period since 2000. Even though labor productivity accelerated another 0.35 percentage point, the growth contributions from IT capital deepening and MFP advances in IT-producing sectors dropped back substantially. At the same time, MFP growth strengthened in the rest of nonfarm business, adding roughly ¾ percentage

---

32. Jorgenson (2001) argues that the steeper declines in semiconductor prices reflected a shift from three-year to two-year technology cycles starting in the mid-1990s. Aizcorbe, Oliner, and Sichel (2006) report that shorter technology cycles drove semiconductor prices down more rapidly after 1995, but they also estimated that price-cost markups for semiconductor producers narrowed from 1995 to 2001. Accordingly, the faster price declines in the late 1990s—and the associated pickup in MFP growth—partly reflected true improvements in technology and partly changes in markups. These results suggest some caution in interpreting price-based swings in MFP growth as a proxy for corresponding swings in the pace of technological advance.

point to annual labor productivity growth during 2000-06 from its 1995-2000 average. And, given the minimal growth in hours worked after 2000, even the anemic advance in investment outlays led to a positive swing in the growth contribution from non-IT capital deepening ("Other tangible capital").[33]

All in all, table 1 indicates that IT-related factors retreated from center stage after 2000 and that other factors—most notably, a surge in MFP growth outside the IT-producing sectors—were responsible for the continued rapid advance in labor productivity as reported in the published data.[34] Nonetheless, for the entire period since 1995, the use and production of IT capital are important, accounting for roughly two-thirds of the post-1995 step-up in labor productivity growth. The next section of the paper examines whether the inclusion of intangible capital changes this characterization.

We conclude this discussion with two points. The first concerns the use of the year 2000 as the breakpoint for comparing the boom period of the late 1990s with more recent years. We chose 2000 rather than 2001 to avoid splitting the two periods at a recession year, which would have accentuated the need for cyclical adjustments. However, our main findings are robust to breaking the two periods at 2001. Second, our big-picture results are very similar to those in Jorgenson, Ho, and Stiroh,[35] which contains the latest estimates from the framework pioneered by Dale Jorgenson. Consistent with our findings, their framework emphasizes the role of IT in explaining the step-up in labor productivity growth during 1995-2000. It also shows a reduced contribution from IT after 2000, which was more than offset by other factors. The differences in

33. The combined effect of adjustment costs and factor utilization remained essentially zero after 2000. Although the deceleration in investment spending after 2000 eliminated the negative effect of adjustment costs, the net decline in the workweek pushed the utilization effect into negative territory.
34. Of course, MFP growth is a residual, so this result speaks only to the proximate sources of growth and does not shed light on the more fundamental forces driving MFP growth.
35. Jorgenson, Ho, and Stiroh (2007).

results are relatively minor and largely stem from the broader sectoral coverage in the Jorgenson, Ho, and Stiroh framework. In particular, their framework incorporates the flow of services from owner-occupied housing and consumer durable goods into both output and capital input. The stocks of these assets have grown rapidly since the mid-1990s, and so Jorgenson, Ho, and Stiroh's estimates of non-IT capital deepening are larger than those reported here.

Aggregate Growth Accounting with Intangible Capital

The growth accounting analysis in the previous section relies on published data, which exclude virtually all types of intangible capital except software. As argued by Corrado, Hulten, and Sichel,[36] any intangible asset that generates a service flow beyond the current period should be included in the capital stock, and the production of such assets should be included in current-period output. Applying this standard, in their 2006 paper (henceforth CHS) Corrado, Hulten, and Sichel estimated that the intangible investment excluded from the NIPAs amounted to roughly $1 trillion annually over 2000-03, an amount nearly equal to outlays for business fixed investment included in the national accounts, and they constructed a growth accounting system that includes a broad set of intangibles through 2003.

Of total business investment in intangibles, CHS estimate that scientific and nonscientific R&D each accounted for about 19 percent during 2000-03; computerized information, which is comprised mostly of the software category already included in the NIPAs, accounted for 14 percent; brand equity accounted for 13 percent; and firm-specific organizational capital accounted for about 35 percent. The last category contains many well-known examples of the successful deployment of intangible capital, including Wal-Mart's supply-chain technology,

36. Corrado, Hulten, and Sichel (2005, 2006).

Dell's build-to-order business model, and Intel's expertise in organizing semiconductor production.[37]

The CHS estimates of intangible investment and capital are a valuable addition to the literature, but the source data for their series are currently available only through 2004 or 2005. Thus their approach cannot be used to develop growth accounting estimates that are as timely as those based on published data. As an alternative, we construct a data system for intangibles that runs through 2006, based on the framework in BFOS. In the BFOS model, firms use intangible capital as a complement to their IT capital. Because of this connection to IT capital, we can generate estimates of intangible investment and capital from published data on IT capital and related series.

BFOS used their model for a more limited purpose: to specify and estimate regressions to discern whether intangibles could explain the MFP growth patterns in published industry data. They did not formally build intangibles into an integrated growth accounting framework along the lines of CHS. That is precisely what we do here.[38]

Description of the Model

The basic features of the BFOS model are as follows. Firms have a (value-added) production function in which IT capital and intangible capital are complementary inputs:

$$V_t = F\left[G\left(K_t^{IT}, R_t\right), K_t^{NT}, L_t, z_t\right], \tag{8}$$

where K_t^{IT}, R_t, and K_t^{NT} denote IT capital, intangible capital, and tangible capital other than IT capital, respectively; L_t is labor input; and z_t is the level of technology. For simplicity, BFOS

37. See Brynjolfsson and Hitt (2000), Brynjolfsson, Hitt, and Yang (2002), and McKinsey Global Institute (2002) for interesting case studies regarding the creation of organizational capital.
38. The BFOS model focuses on intangibles that are related to information technology. This is a narrower purview than in Corrado, Hulten, and Sichel (2005, 2006), who develop estimates for a full range of intangible assets, regardless of their connection to IT. Although we do not provide a comprehensive accounting for intangibles, we highlight the intangible assets that are central to an assessment of the contribution of information technology to economic growth.

assume that there are no adjustment costs and that factor utilization does not vary. The function G that combines IT capital and intangible capital is assumed to take the constant elasticity of substitution form:

$$(9) \qquad G\left(K_t^{IT}, R_t\right) = \left[a\left(K_t^{IT}\right)^{(\sigma-1)/\sigma} + (1-a)\left(R_t\right)^{(\sigma-1)/\sigma}\right]^{\sigma/(\sigma-1)}$$

where σ is the elasticity of substitution between K_t^{IT} and R_t, and a governs the income share of each type of capital.

Because K_t^{IT} and R_t are separable from other inputs, firms minimize costs by first choosing the optimal combination of K_t^{IT} and R_t and then selecting other inputs conditional on this choice. For the first-stage optimization, the usual first-order condition sets the ratio of the marginal products of K_t^{IT} and R_t equal to the ratio of their user costs, which implies

$$(10) \qquad R_t = K_t^{IT}\left(\frac{1-a}{a}\right)^{\sigma}\left(\frac{r_t^{IT}}{r_t^{R}}\right)^{\sigma},$$

where r_t^{IT} and r_t^{R} denote the respective user costs. Equation 10 implies the following expression for the growth of intangible capital:

$$(11) \quad \dot{R}_t = \dot{K}_t^{IT} + \sigma\left(\dot{r}_t^{IT} - \dot{r}_t^{R}\right).$$

Importantly, equation 11 enables us to calculate a model-implied series for the growth rate of intangible capital based solely on data for IT capital and user costs and on an assumed value for the elasticity of substitution between intangible capital and IT capital. No direct data on intangible capital are required. We chain together the time series of growth rates from equation 11 to produce an indexed series for the level of real intangible capital, R.

To implement equation 11, we calculate $\dot{K}_t^{IT}$ and $\dot{r}_t^{IT}$ from the same BLS data that we used in the previous section. We also need to specify the user cost for intangible capital (r_t^R) and the elasticity of substitution between IT capital and intangible capital (σ). We use data from CHS to calculate $\dot{r}_t^R$ and σ, as described next.

CHS measure the user cost of intangible capital in accord with the standard Hall and Jorgenson formulation:[39]

$$r^R = p^R\left(\rho+\delta^R-\Pi^R\right)T^R, \tag{12}$$

where p^R is the price index for this type of capital, ρ is the nominal rate of return net of depreciation, δ^R is the depreciation rate, Π^R is the expected capital gain over and above that captured in the depreciation rate, and T^R accounts for the tax treatment of intangible assets. Equation 12 is identical to the user cost formula that we employ for all other types of capital in our growth accounting framework. We adopt CHS's specification of each term in the user cost formula.

To select a value for the elasticity of substitution σ, we examined the CHS series for the income shares of IT capital and intangible capital.[40] If σ were equal to one (the Cobb-Douglas case), the ratio of the IT income share to the intangible income share drawn from data in CHS (which we denote by $\alpha_t^{R,CHS}$) would be constant. In fact, the ratio of the IT income share to the intangible income share trends upward in the CHS data. Given that the user cost of IT capital has fallen relative to that of intangible capital, the upward trend in the share ratio implies more substitution toward IT capital than would occur in the Cobb-Douglas case. We find that setting σ to 1.25 approximates the upward trend in the share ratio.

To complete the system, we need a nominal anchor to convert the indexed series for R_t to

39. Hall and Jorgenson (1967).

40. Specifically, for the income share of intangible capital, we use the income share series for "New CHS intangibles," that is, those intangibles over and above those included in the NIPAs. We then adjust this series downward to account for the fact that some CHS intangibles are not related to IT and thus do not fit in the BFOS framework. As a crude adjustment, we remove the income share associated with brand equity and one-third of the income share for other components of "New CHS intangibles."

dollar values. For the nominal anchor, we require that the average income share of intangible capital in our framework over 1973-2003 (denoted $\bar{\alpha}_t^{R,BFOS}$) equal the average value of the CHS-based share over the same period:[41]

(13) $\bar{\alpha}^{R,BFOS} = \bar{\alpha}^{R,CHS}$.

To satisfy equation 13, we scale the indexed levels series for intangible capital, R_t, by ψ. The income share for intangible capital in year t is then

(14) $$\alpha_t^{R,BFOS} = \frac{r_t^R R_t \psi}{p_t V_t + r_t^R R_t \psi},$$

where the denominator equals the sum of published nonfarm business income and the income accruing to intangible capital. We average equation 14 over the period 1973-2003, substitute the average share into the left-hand side of equation 13, and solve for the scaling factor ψ. We then apply this scaling factor to the indexed levels series for R_t and denote the resulting series for real intangible capital by R_t^*. Given R_t^*, the associated series for real intangible investment comes from the standard perpetual inventory equation:

(15) $N_t^* = R_t^* - \left(1 - \delta^R\right) R_{t-1}^*$.

We calculate growth in real intangible investment from the series for N_t^*.

We now have all the pieces we need to incorporate intangibles into our growth accounting framework. An important point is that including intangible assets affects both the output and the input sides of the production accounts. On the output side, the growth of production equals a weighted average of growth in real intangible investment $\dot{N}^*$ and growth in published real nonfarm business output. The weight for each component equals its share in the augmented measure of current-dollar output. On the input side, the total contribution from capital

41. We use 2003 as the final year for this calculation because that is the last year of data in CHS.

now includes a term for intangible capital, calculated as the income share for intangible capital times the growth rate of this capital in real terms, $\alpha_t^{R,BFOS} \times \dot{R}^*$. The income shares for all other inputs are scaled down so that the shares (including that for intangible capital) sum to one.[42]

Results

The results from this augmented growth accounting framework, shown in table 2, differ in important respects from the results based on published data. As can be seen by comparing the first two lines, labor productivity growth during 1995-2000 becomes stronger once we include intangibles, but it becomes less robust during 2000-06. Indeed, in the augmented framework, the productivity advance since 2000 is estimated to be well below that posted during 1995-2000, reversing the relative growth rates for the two periods based on published data. This reversal arises from the time profile for real investment in intangibles. As shown in the lower part of the table, real intangible investment is estimated to have surged during 1995-2000, boosting growth in aggregate output, and then retreated during 2000-06.

The growth contribution from intangible capital deepening ("New intangible capital" in table 2) follows the general pattern for IT capital, moving higher during 1995-2000 and then falling back. This similarity reflects the explicit link between intangible capital and IT capital in the BFOS model. The lower part of the table provides full detail on the growth of intangible capital and its determinants from equation 11. Despite the broadly similar growth contour for intangible capital and IT capital across periods, intangible capital increases much less rapidly than IT capital in each period, because of the quality-adjusted declines in IT prices that cause the user cost of IT capital to trend lower. This user cost effect became more pronounced during

42. See Yang and Brynjolfsson (2001) for an alternative approach to incorporating intangibles into a standard growth accounting framework. Their approach relies on financial market valuations to infer the amount of unmeasured intangible investment and shows that, through 1999, the inclusion of intangibles had potentially sizable effects on the measured growth of MFP.

1995-2000—when the prices for IT capital goods fell especially rapidly—restraining the growth of intangible capital even though the growth of IT capital jumped.

Taken together, the revisions to the output and the input sides of the growth accounting equation imply a revised path for MFP growth, after controlling for the effects of adjustment costs and factor utilization ("Growth of MFP excl. above effects"). The inclusion of intangibles leaves a somewhat smaller imprint on MFP growth than on the growth of labor productivity, as the revisions to the two sides of the growth accounting equation are partly offsetting. Consistent with the more muted revision from the published data, the path for MFP continues to show the fastest growth after 2000. However, the pickup in MFP growth from 1995-2000 to 2000-06, at 0.04 percentage point, is negligible compared with that indicated by published data (see the equivalent line in table 1).

Robustness Checks

The BFOS model imposes a strictly contemporaneous relationship between the growth of intangible capital and the growth of IT capital. This relationship may be too tight, as the two forms of capital accumulation may be subject to (unmodeled) adjustment costs and differences in project length from the planning stage to final rollout.

To examine the robustness of our results, we consider alternative timing assumptions for the growth of intangible capital. The first two alternatives smooth the growth of intangible capital without introducing leads or lags relative to the growth in IT capital. The idea is that some projects to produce intangible capital may be long-lived and thus may not display the same stops and starts as purchases of IT capital. We implement this timing change by using a three-year or a five-year centered moving average for the growth rate of IT capital and its user cost on the right-hand side of equation 11. The third timing change allows intangible capital growth to lag IT capital growth by a year but does not affect the relative volatility of the series. This timing assumption embeds the often-expressed view that firms take time to accumulate the intangible capital needed to fully leverage their IT investments.

Our reading of the literature suggests that the first two alternatives fit the facts better than the introduction of a systematic lag from IT capital to intangible capital. Case studies published elsewhere portray the installation of IT capital and associated changes in business practices and organization as being interwoven rather than strictly sequential.[43] Sinan Aral, Erik Brynjolfsson, and D. J. Wu support this view, noting that "… [as] firms successfully implement IT (and complementary intangible investments) and experience greater marginal benefits from IT investments, they react by investing in more IT," a process they characterize as a "virtuous cycle."[44] Nonetheless, we consider the scenario with the lagged accumulation of intangible capital for the sake of completeness.

As the top panel of table 3 shows, these alternative timing assumptions have some effect on the period-by-period growth of real intangible capital but do not change the basic result, namely, that this type of capital essentially has not grown since 2000. The series for intangible investment, shown in the bottom panel of the table, is also reasonably robust to alternative timing assumptions. In each case, real intangible investment is estimated to have declined since 2000. As a further robustness check, the table also displays the CHS series for intangible capital and intangible investment, which we have extended through 2005 based on some of the key source

43. Brynjolfsson and Hitt (2000), Brynjolfsson, Hitt, and Yang (2002), and McKinsey Global Institute (2002).
44. Aral, Brynjolfsson, and Wu (2006, p. 2). Some interpret the econometric results in Brynjolfsson and Hitt (2003) as support for a lag between the installation of IT capital and the accumulation of complementary capital. We believe this interpretation is incorrect. Brynjolfsson and Hitt show that the firm-level effect of computerization on MFP growth is much stronger when evaluated over multiyear periods than when evaluated on a year-by-year basis. Importantly, however, the variables in their regression are all measured contemporaneously, whether over single-year or multiyear periods. Accordingly, their results suggest that the correlation between the growth of IT capital and intangible capital may be low on a year-by-year basis, but that a stronger *contemporaneous* correlation holds over longer periods, boosting the measured effect on MFP growth.

data in their framework. (This is a preliminary extension of the CHS series for illustrative purposes only and should not be regarded as official CHS data.) The extended CHS series for intangible investment and capital exhibit patterns across periods that are broadly similar to those in our series. Notably, the CHS series decelerate sharply after 2000, and the growth rates for 2000-05 are the weakest for the three periods shown, confirming an important qualitative feature of our estimates. Because the CHS series are constructed independently from the series in this paper, the qualitative correspondence between them lends credibility to the basic thrust of our results, if not to the precise figures.

Table 4 explores the growth accounting implications of the alternative timing assumptions for intangible capital. For each timing assumption we show three key variables: growth in labor productivity, the growth contribution from intangible capital deepening, and MFP growth (after controlling for the effects of adjustment costs and factor utilization). Most features of the baseline results are robust to the alternative assumptions. In every case, labor productivity is estimated to have grown more rapidly during 1995-2000 than during 2000-06, reversing the relative growth rates based on published data. In addition, the growth contribution from intangible capital deepening is always largest during 1995-2000 and then drops back to essentially zero during 2000-06. Finally, although the alternative timing assumptions generate a larger step-up in MFP growth after 2000 than in the baseline, they nonetheless temper the increase by 0.2 to 0.3 percentage point relative to the published data.

Industry-Level Productivity

We now turn to the industry origins of U.S. productivity growth during the late 1990s and after 2000. The aggregate data show that the sources of productivity growth changed after 2000, which suggests that the industry-level origins of aggregate productivity growth and the underlying forces may have also changed. To explore this, we construct productivity accounts for sixty industries that span the U.S. private economy from 1988 to 2005. Although measurement error, omitted variables, and endogeneity problems always make it difficult to identify the sources of productivity gains, we make some progress by exploiting cross-sectional variation in industry productivity over time and by examining the link between productivity and observable factors such as IT intensity and changing profit shares.

The industry analysis presented here focuses on labor productivity, reflecting our interest in understanding the industry origins of aggregate labor productivity growth. Moreover, we do not have the detailed data on labor quality, intangible investment, or adjustment costs at the industry level necessary to create comparable estimates of MFP growth. To the extent that intangible capital is correlated with IT investment, however, one can interpret the IT intensity results as broadly indicative of the whole suite of activities that are complementary to IT.

Output Measures, Data, and Summary Statistics

Output Measures. Industry output can be measured using either a gross output or a value-added concept, each with its advantages and disadvantages.[45] Gross output corresponds closely to the conventional idea of output or sales and reflects all inputs including capital, labor, and intermediate energy, materials, and services. Value added, by contrast, is a somewhat artificial concept that strips out the contribution of intermediate inputs and incorporates only capital and labor.

Although both value added and gross output are used for productivity analysis, we favor gross output. Empirical work by, among others, Michael Bruno; J. R. Norsworthy and David Malmquist; Jorgenson, Frank Gollop, and Barbara Fraumeni rejects the existence of value-added

45. For background on industry productivity analysis, see Jorgenson, Gollop, and Fraumeni (1987), Basu and Fernald (1995, 1997, 2001, 2007), Nordhaus (2002b), Stiroh (2002a, 2002b), Triplett and Bosworth (2004), and Bosworth and Triplett (2007).

functions on separability grounds.[46] Basu and Fernald show that using value-added data leads to biased estimates and incorrect inferences about production parameters.[47] A later contribution by the same authors argues against the value-added function because failure of the neoclassical assumption about perfect competition implies that some of the contribution of intermediate inputs remains in measured value-added growth.[48] Value added has the advantage, however, that it aggregates directly to GDP.

Data. We use three pieces of U.S. industry-level data—output, hours, and capital stock—from government sources. The first two create a panel of average labor productivity (ALP) across U.S. industries, and the third is used to develop measures of the intensity of the use of IT. One practical difficulty is the recent conversion of the industry data from the Standard Industrial Classification (SIC) system to the North American Industrial Classification (NAICS) system, which makes it difficult to construct long historical time series or to directly compare the most recent data with earlier results.

BEA publishes annual data on value added and gross output for sixty-five industries that together make up the private U.S. economy.[49] These data, which are based on an integrated set of input-output and industry production accounts, span 1947-2005 for real value added and 1987-2005 for real gross output. Although BEA also publishes various measures of employment by industry, it does not provide industry-level series on hours worked. We obtain hours by industry from the Output and Employment database maintained by the Office of Occupational Statistics and Employment Projections at BLS. Complete data on total hours for all industries begin in

46. Bruno (1978); Norsworthy and Malmquist (1983); Jorgenson, Gollop, and Fraumeni (1987).
47. Basu and Fernald (1995, 1997).
48. Basu and Fernald (2001).
49. Howells, Barefoot, and Lindberg (2006).

1988.[50] Because these hours data are currently available only to 2004, we use the growth rate of full-time equivalent employees for the disaggregated industries, from BEA data, to proxy for hours growth in 2005.

We create two measures of industry ALP—real value added per hour worked and real gross output per hour worked—by combining the BEA output data with the BLS hours data across industries for 1988 to 2005.

The third data source is the Fixed Asset accounts from BEA for nonresidential capital. These data include forty-six different types of nonresidential capital for sixty-three disaggregated NAICS industries since 1987. To estimate capital services we map the asset-specific service prices from Jorgenson, Ho, and Stiroh onto these assets and employ Tornqvist aggregation using the service price and a two-period average of the capital stock for each asset in each industry.[51] The resulting measure of capital services is an approximation, because we miss industry variation in rates of return, asset-specific inflation, and tax code parameters. Nevertheless, it captures the relatively high service prices for short-lived assets such as IT capital, defined as above to include computer hardware, software, and communications equipment.

We combine these three sources of data to form a panel from 1988 to 2005 for a private industry aggregate, fifteen broad sectors, and sixty disaggregated industries. The fifteen-sector breakdown follows BEA's convention, except that manufacturing is broken into durables and nondurables. The number of disaggregated industries is smaller than that available from either BEA or BLS, because of the need to generate consistently defined industries across all data sources. All aggregation is done via Tornqvist indices, except for hours, which are simply

---

50. The underlying sources of these data are the BLS Current Employment Survey (for wage and salary jobs and average weekly hours), the Current Population Survey (for self-employed and unpaid workers, agricultural workers, and household employment), and unemployment insurance tax records.
51. Jorgenson, Ho, and Stiroh (2007).

summed. Both the broad sectors and the disaggregated industries sum to the private industry aggregates of nominal output from BEA, hours from BLS, and nominal nonresidential capital from BEA. The list of industries and their 2005 value added are reported in appendix table A-1.

Summary Statistics. Table 5 reports estimates of labor productivity growth from our industry data and compares them with the latest estimates from BLS. The first two lines of the top panel report average annual growth of ALP for the BLS business and nonfarm business sectors, and the third line reports the private industry aggregate described above. Although our private industry aggregate grows somewhat less rapidly than the BLS aggregates, all three series show similar trends: a pickup of ALP growth of about 1 percentage point after 1995 and a smaller increase after 2000.

The second panel of table 5 reports estimates for the fifteen broad NAICS sectors. These sectors range in size from the very large finance, insurance, real estate, rental, and leasing sector, at 23.3 percent of 2005 value added, to the very small agriculture, forestry, fishing, and hunting sector, at only 1.1 percent. In terms of ALP growth, eight of these fifteen sectors, which accounted for 73 percent of value added in 2005, showed faster productivity growth over 1995-2000 than over 1988-95.[52] The further acceleration in aggregate productivity after 2000 occurred in seven sectors, which accounted for only 44 percent of 2005 value added. Although productivity in the large retail trade, wholesale trade, and finance sectors all decelerated after 2000, the two trade sectors continued to post impressive productivity gains through 2005.

The pickup in aggregate productivity growth in the mid-1990s appears to have originated in different sectors than did the subsequent step-up in 2000. Six sectors (agriculture, durable goods, wholesale trade, retail trade, finance, and arts and entertainment) show an acceleration

52. As a comparison, Stiroh (2001, 2002b) reported an acceleration of ALP after 1995 for six of ten broad sectors, which accounted for the majority of output using earlier vintages of SIC data.

after 1995 but a deceleration after 2000, whereas five sectors (construction, nondurables, utilities, information, and other services) show the opposite pattern. Together these eleven sectors produced 72 percent of value added in 2005. In their analysis of MFP growth, Corrado and others reach a similar conclusion, although Bosworth and Triplett emphasize the continued importance of service industries as a source of aggregate productivity growth.[53]

Table 5 also summarizes, in the third and fourth panels, the disaggregated industry data by reporting the mean, median, and hours-weighted mean productivity growth rates across these industries for gross output and value added, respectively. One interesting observation is the divergence in trends between gross output and value-added measures of productivity: the post-1995 gains are strongest for gross output, whereas the post-2000 gains are strongest for value added. Both series incorporate the same hours data, so that this divergence directly reflects differences between the gross output and value-added output measures.

It is beyond the scope of this paper to investigate this divergence further. For completeness, we report results for both gross output and value added, although, again, we prefer gross output because it is a more fundamental measure of production and does not require additional assumptions about the nature of the production function.

Finally, we emphasize that there is enormous heterogeneity among the disaggregated industries that lie beneath these summary statistics, both within time periods and across time. For example, thirty-seven of the sixty industries, which accounted for nearly 60 percent of aggregate output, experienced an acceleration of productivity after 1995 but a decline after 2000, or vice versa. This highlights the widespread churning and reallocation of resources among industries, which we show to be an important source of aggregate productivity gains.

53. Corrado and others (2007); Bosworth and Triplett (2007).

Industry Origins of the Aggregate Productivity Gains

We now review how the data for the disaggregated industries can be aggregated to form economywide productivity estimates, and we employ this familiar framework to identify the industry origins of the aggregate productivity gains over 1988-2005.

Decomposition and Reallocations. At the industry level, real value added is defined implicitly from a gross output production function as

$$(16) \qquad \dot{Y}_i = \alpha_i^V \dot{V}_i + (1-\alpha_i^V)\dot{M}_i,$$

where α_i^V is the average share of nominal value added in nominal gross output for industry *i,* and M_i denotes real intermediate inputs.[54] One attractive property of industry value added is that it aggregates to a simple expression for growth in aggregate value added:

$$(17) \qquad \dot{V} = \sum_i v_i \dot{V}_i,$$

where v_i is the average share of industry *i*'s nominal value added in aggregate nominal value added. Aggregate hours worked, *H*, is the simple sum of industry hours, H_i,

$$(18) \qquad H = \sum_i H_i,$$

and aggregate labor productivity is defined as $ALP^V = V/H$.

Equations 16, 17, and 18 can be combined to yield the following decomposition of ALP growth:[55]

54. BEA uses the "double deflation" method to estimate real value added for all industries as the difference between real gross output and real intermediate inputs (Howells, Barefoot, and Lindberg, 2006). Basu and Fernald (2001) show that this can be approximated, as in equation 16, by defining gross output as a weighted average of value added and intermediate input growth.
55. As in Stiroh (2002b).