ROBERT J. GORDON
Northwestern University

Exploding Productivity Growth: Context, Causes, and Implications

WILL POTENTIAL OUTPUT GROW in the future at a 4 percent annual rate, as several of the more optimistic business economists assume, or at the pathetic 1.8 percent annual rate assumed into the distant future by the trustees of the Social Security Administration?[1] Put differently, will real GDP in seventy-five years be 20 times its current level or a mere 3½ times? Academic research on future supply-side issues has focused mainly on the causes of the post-1995 productivity growth revival, but the growth rate of potential output is of independent interest. Variations in four factors—population growth, labor force participation, the unemployment rate, and hours worked per employee—can create significant differences between the long-run path of potential output and that of trend productivity growth. These differences matter for numerous issues,

1. A long-term real GDP growth path of 4 percent a year and a 4 percent permanent unemployment rate are endorsed in Glassman (2002) and displayed in his figure 4. The long-run real GDP growth path in the Social Security projections tapers down from 3.0 percent to 2.0 percent between 2003 and 2015, rests at 1.9 percent from 2020 to 2040, and then falls to 1.8 percent between 2040 and 2080 (Board of Trustees, 2003, table V.B.2, p. 99).

including long-run fiscal policy, the solvency of entitlement programs, the balance of world saving and investment, and the role of the United States as an engine of growth for the rest of the world.

This paper has three goals. The first is to forecast growth in U.S. potential real GDP, not for the full seventy-five-year horizon of the Social Security trustees, but for the more modest but still daunting span of the next two decades. It brings together recent research both about productivity and about the likely future behavior of the other four factors, especially population growth, that matter for potential output growth. The need to predict future population growth in turn requires an exploration of the determinants of trends in fertility and mortality rates, as well as the likely future trend of net immigration into the United States.

The second goal of the paper, connected closely with the first, is to interpret the extraordinary productivity performance of the United States since 1995 and especially since mid-2000. Far from slowing in response to the 2001 recession and the collapse of investment in information and communications technology (ICT) after mid-2000, growth in labor productivity actually accelerated from an average of 2.56 percent a year between 1995:4 and 2000:2 to 3.46 percent a year between 2000:2 and 2003:2. Should a forecast of future productivity growth use as its precedent the average behavior of actual productivity growth over the past two years, the past eight years, or some longer interval?

The third goal of the paper, related to the first two, is to provide a new breakdown of past U.S. economic growth into its trend and cyclical components. In assessing long-term growth performance over some historical period, one would not want to include the portion of real GDP growth contributed by a sharp difference in cyclical conditions, for example between the 7.6 percent unemployment rate of mid-1992 and the 4.0 percent rate of early 2000. This paper bases its cyclical analysis on an identity that links real GDP to productivity, the employment rate, and several other variables. This analysis uncovers important changes in cyclical behavior between the earlier postwar downturns and the two recent jobless recessions and recoveries (1990–93 and 2001–03). One particularly important difference is the strength of productivity growth and the weakness of payroll employment growth in both of the most recent episodes and especially the latest.

To predict the future, one naturally starts with the past. But how much of the past? This is not a question on which either economic or statistical

theory provides much guidance. Productivity growth was much faster after 1995 than between 1972 and 1995, but if one is trying to make forecasts twenty or more years into the future, the slow-growth period has at least some relevance, as do the more rapid growth periods 1950–72 and 1995–2003. Sometimes one has good economic reasons for looking at only part of history, believing strongly, for instance, that the economic dislocations of the Great Depression and World War II will not recur. When making demographic forecasts of future population growth, it is quite certain that fertility rates will never return to those of the baby-boom era of 1947 to 1963, when the average female of childbearing age had between three and four children.[2] But is it possible to be so sure which decades into the past are relevant for predictions of immigration?

This paper of necessity is limited to data for the United States. But one should keep in mind the dimension of "American exceptionalism" that pervades international comparisons. In forecasting over a period as long as two decades, it is important to at least consider the possibility that the United States could become more like other developed countries in some aspects in which it has long been unusual, instead of continuing to diverge. The most important aspects of American exceptionalism in the recent past have been the productivity growth revival, the absence of a decline in hours of work per employee, the relatively high level of fertility, and the important role of immigration in maintaining population growth.

In examining the third topic, the separation of cycle and trend, I adopt a modern version of Okun's Law and define potential real GDP as the output that the economy can produce at its nonaccelerating-inflation rate of unemployment (NAIRU), also known as the natural rate of unemployment.[3] My analysis of cycles uses the output identity that links potential output to the employment rate and other variables. By estimating regression equations that explain changes in the employment rate and the other variables by changes in detrended real GDP, one can achieve a complete decomposition of each component of the identity into measures of trend, regular cyclical patterns, and residual terms that reveal unusual behavior not predicted by the regressions.

2. The U.S. total fertility rate reached 3.7 births per woman of childbearing age in the interval 1958–61. See Technical Panel on Assumptions and Methods (2003, figure 1, p. 23).
    3. Okun (1962).

This three-way breakdown among trend, cycle, and residuals allows one to address a wide set of issues, including some unusual aspects of the economic boom of 1995–2000, the 2001 recession, and the recovery since late 2001, viewed both on their own and in comparison with past business cycles. For instance, why has payroll employment been so much weaker in 2002–03 than in 1991–92, the episode that launched the term "jobless recovery," even though the unemployment rate is still well below its 1992 peak? Is this discrepancy partly due to a greater than normal downturn in the labor force participation rate (LFPR), suggesting a rise in discouraged workers? Have hours per employee held up better than usual, so that firms are hiring less because they are working their existing employees harder? Or is there a measurement issue in the form of unusual behavior in the ratio of employment as reported from the Bureau of Labor Statistics (BLS) payroll survey to that reported from the BLS household survey (the payroll-to-household employment ratio)?

The paper tackles the three goals in reverse order, beginning with the output identity and the econometric decomposition of trend and cycle. Using those results, the paper proceeds to discuss unusual aspects of productivity and employment behavior over the period 2000–03. Finally, armed with insights from these analyses, the paper proceeds to the two-decade long-term forecasts.

The paper begins by establishing notation and defining the output identity. The first examination of the data over the 1954–2003 period is based on defining log-linear trends through certain designated benchmark quarters (hereafter called the trends-through-benchmark, or TTB method) and examining the behavior of those trends for each component of the output identity.[4] The next step is to go beyond the TTB detrending method with two alternative statistical techniques for estimating trends, the Hodrick-Prescott filter and the Kalman filter. A careful consideration of the differences and limitations of all three techniques leads to a hybrid concept of the potential output trend, based on a mix of all three methods.

Having established trends for each variable, I then reassess, for each component of the identity, the Okun's Law coefficient using both a non-

---

4. Two of my previous Brookings Papers (Gordon, 1979, 1993) were limited to an examination of the trend and cyclical behavior of labor productivity. An additional paper (Gordon, 1984) explored the trend and cyclical behavior of the output identity but primarily limited its attention to the 1981–82 recession and the early phase of the 1983–84 recovery. Earlier uses of the output identity were by Clark (1983) and Woodham (1984).

statistical examination of peak and trough ratios of output to trend and a regression analysis. The regression analysis yields a distinction between normal cyclical behavior and residual, unexplained behavior that emphasizes the differences between the two jobless recoveries of 1991–92 and 2001–03 and previous cyclical episodes.

The econometric results, which reveal a substantial residual component in productivity and aggregate hours in 2001–03, raise both measurement and substantive issues. A central puzzle is the interpretation of rapid productivity growth after 2000 in the context of the crash of ICT investment along with the stock market. Was the role of ICT investment in the 1995–2000 productivity revival exaggerated, and if so, why? Balancing the overemphasis on ICT investment in the recent literature has been a neglect of other aspects of the post-2000 acceleration, including the delayed benefits of intangible capital, the intensity of corporate cost cutting, and increased outsourcing to foreign countries.

The paper concludes by discussing some issues involved in projecting the components of the output identity two decades into the future, with an emphasis on productivity, population growth, and changes in hours per employee. This section contains no econometrics, but rather introduces the kinds of decisions that must be made in carrying out such forecasts, as well as a set of forecast numbers intended to focus subsequent discussion rather than to provide precise and definitive answers.

Trends, Cycles, and Residuals

This section establishes the notation to be used throughout the paper and defines the output identity. It then explores which components of the output identity account for differences between the growth rates of output per capita and of productivity over time. Finally, it considers alternative methods of distinguishing trends from cyclical fluctuations in output and its components.

The Output Identity: Notation and Definitions

The output identity is equally useful for distinguishing the components of long-run growth in potential GDP and for assessing the cyclical behavior of Okun's Law, the regular empirical relationship that allows

one to decompose fluctuations in output away from trend into fluctuations in the employment rate, productivity, the LFPR, and hours per employee. I begin with a simple version of this identity, which decomposes real GDP (Q) into output per hour (Q/A), aggregate hours per employee (A/E), the employment rate (E/L), the LFPR (L/N), and the working-age population (N).[5]

(1)
$$Q \equiv \frac{Q}{A} \cdot \frac{A}{E} \cdot \frac{E}{L} \cdot \frac{L}{N} \cdot N.$$

At this stage I suppress time subscripts, since all of the variables in equation 1 and the subsequent versions of the output identity are contemporaneous. The right-hand side contains four elements that typically display procyclical behavior (albeit with different sets of leads and lags relative to total output), namely, output per hour, hours per employee, the employment rate, and the LFPR. One would expect no response of the working-age population (N) to the business cycle.

The identity in equation 1 is a useful starting point but cannot be used for empirical analysis of cyclical issues, because it ignores the fact that productivity data are published on a quarterly basis only for the private business sector and several subaggregates, particularly the nonfarm private business (NFPB) sector. Also of interest are differences in the behavior of employment as reported by the household survey and by the establishment (or payroll) survey. These complications require that equation 1 be expanded to identify those variables that refer to the NFPB sector (designated here with a superscript B). Variables without superscripts refer to the total economy, except that total employment in the payroll survey is referred to as E^P, to distinguish it from total employment as measured in the household survey (E):

(2)
$$Q \equiv \frac{Q^B}{A^B} \cdot \frac{A^B}{E^B} \cdot \frac{E}{L} \cdot \frac{L}{N} \cdot N \cdot \frac{Q/E^P}{Q^B/E^B} \cdot \frac{E^P}{E}.$$

Equation 2 differs from equation 1 in the last two terms. The next to last term, which I call the "mix effect," measures the ratio of output per payroll employee in the total economy to that in the NFPB sector. The final

---

5. The employment rate E/L is simply 1 minus the unemployment rate, that is $(1 - U/L)$.

term measures the ratio of employment from the payroll survey to that from the household survey.

To simplify the notation and allow for the subsequent treatment of growth rates and ratios of actual to trend, I take the logarithm of each factor in equation 2 and choose a single letter for each. This relationship states that, by definition, the logarithm of real GDP (q) is equal to the sum of the logarithms of the following: output per hour (p), hours per employee in the NFPB sector (h), the employment rate (e), the LFPR (f), the working-age population (n), the mix effect (m), and the payroll-to-household employment ratio (s):

(3) $$q \equiv p + h + e + f + n + m + s.$$

Of interest are the growth rates of the components of equation 3 over longer or shorter intervals, the estimated trends of the components, and the ratios of actual to trend expressed as logarithms.

Using the same notation, I identify trends, ratios of actual to trend, and growth rates in the simplest possible way. The trend of the logarithm of real GDP (q^*) is the sum of the same seven components as in equation 3 when each of the components, say, x, is marked with an asterisk. The log ratio of actual to trend for each component x' also observes the identity in equation 3: $x' = x - x^*$. The regression analysis below will study rates of change in quarterly data. Adding time subscripts, a rate of change is $\Delta x_t = x_t - x_{t-1}$. The dependent variable in the regressions is the rate of change of the deviation of a variable from trend, $\Delta x'_t = \Delta x_t - \Delta x^*_t$. Rates of change in the logarithms of actual values, of trend values, and of deviations from trend all obey equation 3.

Okun's Law can be expressed in this notation either as a relationship among deviations from trend or as a relationship among rates of change. One can define k as the response of the deviations from trend of the employment rate relative to total output:

(4) $$k = e'/q'.$$

The coefficient k was one-third in Okun's original analysis. However, it has long been recognized that unemployment is a lagging indicator, and including both the current and the lagged reaction of unemployment and the employment rate brings the coefficient k up to roughly 0.4 in the new results presented below. If $k = 0.4$, then cyclical movements in q' have

roughly 2.5 times the amplitude of cyclical changes in the employment rate (e'), and the remaining fluctuations in q' by definition must occur among the other components of the output identity.

Growth Rates between Benchmark Quarters

My first use of the output identity is to see which components account for differences between the growth rates of output per capita and of productivity over time. To carry out this analysis, I need to choose intervals that are free of marked business cycle fluctuations. At this initial stage of simple data examination, I choose to calculate growth rates over intervals between certain "benchmark" quarters, which represent roughly the same stage of the business cycle. These quarters are those when the actual unemployment rate is equal to the NAIRU.[6] In each business cycle, starting at the trough, the unemployment rate declines through the NAIRU and in the next recession rises back through the NAIRU; the benchmark quarters occur at the first of these events, the quarter when the actual unemployment rate first declines below the NAIRU. The one exception is the last benchmark quarter: 2001:2 was chosen because at that time the unemployment rate had not fallen through the NAIRU as in earlier episodes such as 1987 and 1994, and instead this quarter was chosen as a time when the actual unemployment rate was rising through the NAIRU. The selection of benchmark quarters is further discussed in appendix A.

Table 1 presents annual growth rates of real GDP and the seven right-hand-side components of the output identity for six intervals between benchmark quarters; thus each of the numbers in table 1 is interpreted as a trend growth rate. Also shown are growth rates for the "last two episodes," that is, the most recent eight quarters of data during 2001–03 compared with eight similar quarters of the 1990–92 period. By definition, as in equation 3, the growth rates of the seven components on the right-hand side of the identity sum to the growth rate of real GDP in the first column.

Of the six periods between 1954 and 2001, the second (1963–72) stands out as having the most rapid growth in real GDP, propelled by rapid productivity and population growth, only partly offset by negative growth in hours per employee. The third period had much slower pro-

---

6. Here the NAIRU takes the values estimated in Gordon (1998), updated in Eller and Gordon (2003).

Table 1. Annual Growth Rate of Real GDP and Components of the Output Identity, 1954–2003

Percent

Period	Real GDP growth	Output per hour[a]	Hours per employee[a]	Employment rate	LFPR	Working-age population	Mix effect[b]	Employment ratio[c]
						Contribution of output identity components		
Between benchmark quarters								
1954:4–1963:3	3.51	2.72	–0.16	–0.02	0.03	1.43	–0.27	–0.22
1963:3–1972:2	4.07	2.54	–0.44	–0.02	0.33	1.80	–0.62	0.49
1972:2–1978:3	3.43	1.72	–0.60	–0.06	0.73	1.94	–0.38	0.08
1978:3–1987:3	2.69	1.20	–0.26	0.00	0.41	1.39	0.06	–0.12
1987:3–1994:4	2.71	1.48	–0.01	0.06	0.23	1.08	–0.42	0.28
1994:4–2001:2	3.21	2.11	–0.31	0.19	0.02	1.24	–0.61	0.57
Last two episodes								
1990:3–1992:3	1.32	1.95	0.01	–1.04	0.10	1.07	–0.26	–0.50
2001:2–2003:2	2.21	4.52	–0.46	–0.92	–0.27	1.20	–1.03	–0.83

Source: See appendix A.
a. Data are for the nonfarm private business sector.
b. Ratio of output per employee in the economy as a whole, from the payroll survey, to that in the nonfarm private business sector.
c. Ratio of employment from the payroll survey to that from the household survey.

ductivity growth but was notable for the fastest growth in the LFPR and in population of any period, partly offset by the fastest rate of decline in hours per employee of any period. After 1978 there was slower growth in productivity, the LFPR, and the population, causing average annual output growth to fall to 2.69 percent in 1978–87 and 2.71 percent in 1987–94. Real GDP growth accelerated by 0.50 percentage point between 1987–94 and 1994–2001, somewhat less than the 0.63-percentage-point acceleration in productivity growth.

When one compares the most recent eight quarters with similar quarters in 1990–92, the biggest difference is that productivity growth has been 2.57 percentage points faster in 2001–03, but real GDP growth has been only 0.89 percentage point faster. This difference is primarily explained by the large negative values for the mix effect and the payroll-to-household employment ratio. Notice that, in all lines but the first in the top panel of the table, the growth rates of the mix effect and the payroll-to-household employment ratio are of opposite sign, whereas in 2001–03 the two growth rates both had relatively large negative signs.

These points are made with a different arrangement of the same data in table 2. There the annual growth rates of real GDP per capita and productivity are shown in the first two columns and their difference in the third.[7] By definition, this difference must equal the sum of the contribution of the three labor market variables (employment rate, the LFPR, and hours per employee) plus the sum of the mix and employment ratio effects. For the six intervals between benchmark quarters, the difference between real GDP per capita and productivity is relatively small. But for the last two episodes the difference is huge, –1.70 percentage points for 1990–92 and –3.51 percentage points for 2001–03. In both periods the explanation is roughly evenly divided between the labor market variables and the mix and employment ratio effects.

Establishing Trends: Choosing a Methodology and a Parameter

To examine the cyclical behavior of components of the output identity, it is necessary to divide actual changes into cyclical and trend elements. And to project the components over the next two decades, the best possible measures of trend behavior, especially in the recent past, are needed.

---

7. Here "per capita" means "per member of the working-age population," not the usual "per member of the total population."

Table 2. Annual Growth Rate of Real GDP and Components of the Output Identity, Alternative Decomposition, 1954–2003

Percent

				Contribution to difference	
Period	Growth in real GDP per capita[a]	Output per hour[b]	Difference[c]	Labor market variables[d]	Mix and employment effects[e]
1954:4–1963:3	2.08	2.72	–0.64	–0.15	–0.49
1963:3–1972:2	2.27	2.54	–0.27	–0.13	–0.13
1972:2–1978:3	1.49	1.72	–0.23	0.07	–0.30
1978:3–1987:3	1.30	1.20	0.10	0.15	–0.06
1987:3–1994:4	1.63	1.48	0.15	0.28	–0.14
1994:4–2001:2	1.97	2.11	–0.14	0.10	–0.04
1990:3–1992:3	0.25	1.95	–1.70	–0.93	–0.76
2001:2–2003:2	1.01	4.52	–3.51	–1.65	–1.86

Source: See appendix A.
a. Per member of the working-age population.
b. Data are for the nonfarm private business sector.
c. Difference between the first and second data columns, equals the sum of the last two data columns.
d. Employment rate, LFPR, and hours per employee.
e. Defined as in table 1.

The growth rates in table 1 represent the TTB method of detrending. Although useful in providing summary information about the growth rates of output and its components over a uniform set of historical intervals, these trends are less useful for the subsequent analysis. When the primary variable of interest is the rate of change of a trend, the TTB method generates instantaneous jumps up or down in this rate of change at the benchmark quarters, as in table 1. As a consequence of these jumps, the TTB method understates trend growth before a positive jump in a trend rate of change and overstates it after such a jump. To create trends that exhibit smooth rather than jumpy behavior in their rates of change, I consider two alternative techniques, the Hodrick-Prescott (H-P) filter and the Kalman filter.

The H-P filter is perhaps the most commonly used detrending method in macroeconomics, presumably because it allows the trend to move continuously and because it is easy to understand and to use in estimation. Its primary flaw is that the estimated trends from any given time series can exhibit wildly different behavior, depending on the smoothness parameter chosen, and that choice of parameter is entirely arbitrary. At one extreme, the choice of a parameter of zero yields a trend that exactly tracks every value of the series being detrended. At the other extreme, a parameter of

infinity yields a single straight log-linear trend. Between zero and infinity, a relatively low value for the smoothness parameter creates a trend series that bends frequently in response to changes in the actual series and hence implies relatively small deviations of the actual values from trend; a high parameter value creates a relatively smooth trend and relatively larger deviations from trend. The other problem is that H-P filtering does not explicitly allow for cyclical variations that are distinct from variations in the trend.

The parameter endorsed by Robert Hodrick and Edward Prescott themselves is a relatively low value (1,600), which implies implausibly large accelerations and decelerations of the trend within each business cycle.[8] For instance, Finn E. Kydland and Prescott use this parameter to conclude that the entire economic boom of the 1960s resulted from an acceleration of trend, rather than a deviation of actual output above trend.[9] This conclusion ignores outside information, such as the fact that the unemployment rate was unusually low and the capacity utilization rate was unusually high.

An alternative technique is the Kalman filter.[10] This can be used to estimate time-varying coefficients in any type of time-series model, whether a complex multiequation model or a single equation. The application used here is even simpler, the estimation of a single time-varying coefficient in a single equation, without allowing the other coefficients in that single equation to change. Detrending methods make little difference for most of the components of the output identity, which tend to evolve smoothly over time, but detrending methods and parameter choices are crucial for

8. There is no justification for this parameter anywhere in the literature. The justification in the original H-P paper (Hodrick and Prescott, 1981) is simply stated as "our prior view." The parameter 1,600 is the square of the ratio of a cyclical deviation from trend to the adjustment per quarter of that trend. In the authors' example, a 5 percent deviation of output from trend would cause the trend growth rate to adjust in the same direction by 1/8 percent per quarter, or $\frac{1}{2}$ percent a year, or by 2 percent a year if that 5 percent output gap were sustained for four years. The value 1,600 is the square of the ratio of the cyclical component (5) to the per-quarter adjustment of trend (1/8): $1,600 = (5/0.125)^2$. Applying this parameter to 1929–33 and assuming an initial trend growth rate of 3 percent a year, a sustained 25 percent output gap would cause the growth rate of potential output to decline from +3 percent a year in 1929 to –7 percent a year in 1933.

9. Kydland and Prescott (1990, chart 2, p. 9).

10. The technique originated in Kalman (1960). A complete (and highly technical) treatment of the filter is contained in Hamilton (1994, chapter 13). As Hamilton shows, the Kalman filter has many uses beyond the estimation of time-varying coefficients.

productivity growth, which has well-recognized cyclical components. Hence my example of the estimation of time-varying coefficients with the Kalman filter explains the change in productivity growth (Δp_t) by a time-varying constant and any set of other explanatory variables (βX_t):

(5) $$\Delta p_t = \alpha_t + \beta X_t + w_t.$$

The next step is to specify a time-series process for the time-varying productivity trend, and the most straightforward is a random walk:

(6) $$\alpha_t = \alpha_{t-1} + v_t.$$

The error terms of this two-equation system are

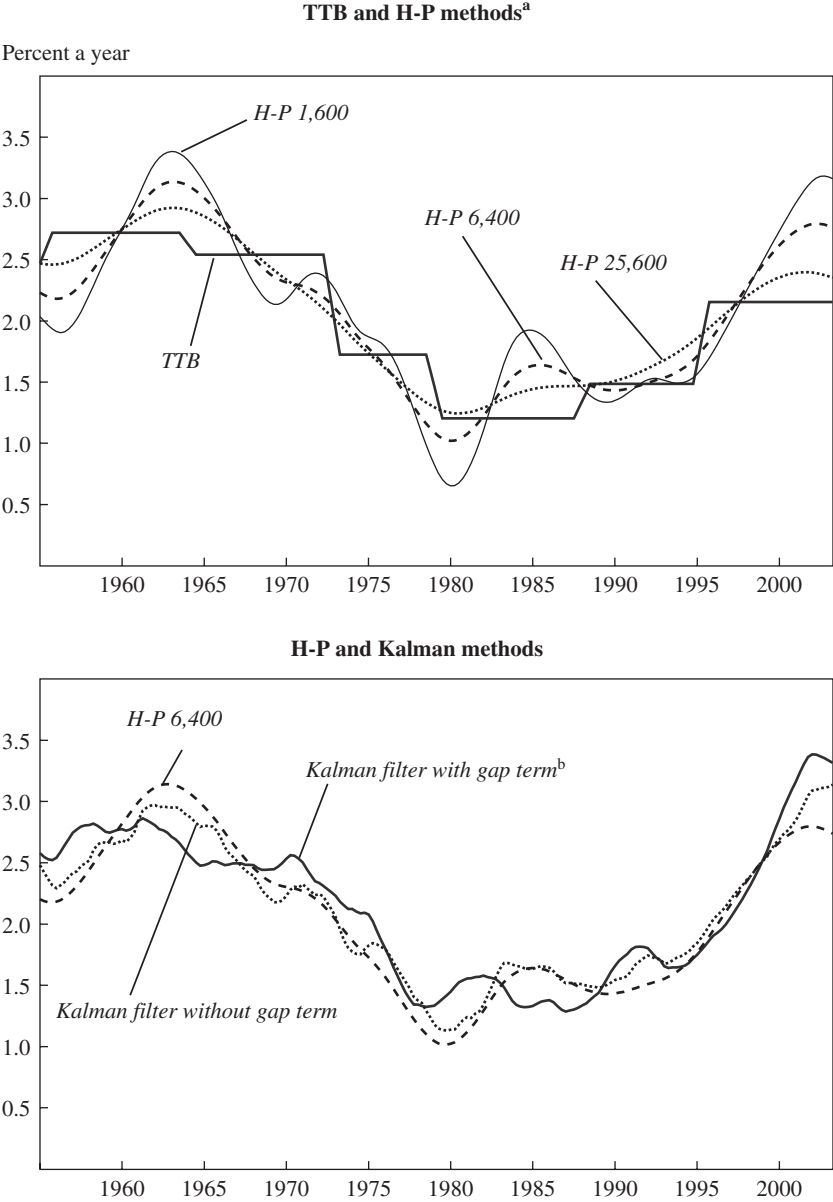$$w_t \sim N(O, \sigma^2); v_t \sim N(O, \tau^2).$$

In estimating this system a smoothness parameter must be specified to control the variance of the random walk process (τ^2), and this then allows a range of trend rates of change to be obtained, from very jumpy to very smooth, just as in the case of the H-P filter.[11]

The two techniques share the weakness that the user must specify the smoothness parameter. The advantage of the Kalman filter is that any additional number of variables (X_t) in equation 5 may be specified to control for determinants of actual productivity changes that do not represent fundamental causes of changes in the trend; for example, the X variables could include changes in unemployment or the output gap, or dislocations caused by short-run events such as strikes or temporary changes in oil prices. In contrast, the H-P filter cannot by its design use any outside information, such as the fact that slow and even negative productivity growth in the early 1980s occurred at a time when the output gap was at a record negative level.

Here I compare alternative measures of the trend growth in productivity in two steps. The top panel of figure 1 compares the TTB method with several alternative H-P parameters, and the bottom panel compares a

---

11. Hamilton (1994, p. 400) provides an exposition in which the movement of the time-varying parameter(s) is governed by an adaptive process in which the current parameter is a weighted average of the lagged parameter and the mean value of the parameter, and the random walk model in equation 6 is a special case where the weight on the lagged parameter is unity. Hamilton cites Doan, Literman, and Sims (1984) as an early example of the use of the adaptive formulation.

Figure 1. Alternative Estimates of Four-Quarter Trend in Growth in Output per Hour, 1954–2003

TTB and H-P methods[a]

Percent a year

H-P and Kalman methods

Source: Author's calculations using data from Bureau of Labor Statistics.
a. H-P, Hodrick-Prescott filter; TTB, trend through benchmark.
b. Cyclical term is the current and four leading values of the change in the output gap (see text).

single H-P trend with two alternative Kalman trends. In the top panel the segments of the line labeled "TTB" plot the same set of productivity growth rates as are shown in the second column of table 1. The TTB trend growth rate is slightly above 2.5 percent a year during 1954–72, slows to 1.2 percent in 1978–87, and then rises to a final value of 2.1 percent in 1994–2001.[12]

The top panel of figure 1 contrasts the TTB trend with H-P trends using three different smoothness parameter values: the 1,600 parameter endorsed by Hodrick and Prescott and used by most H-P users, and two higher (smoother) parameters, 6,400 and 25,600. The excess sensitivity of the 1,600 parameter is evident in the crash dive of trend productivity in response to the back-to-back recessions of 1980–82, its soaring recovery in 1983–86, and then its milder dive during the 1990–91 recession. The 6,400 and 25,600 parameters present successively smoother versions of history, and the higher parameter exhibits little response to low economic output in 1980–82 and none to the 1990–91 recession. However, the 25,600 parameter has the disadvantage that its inability to "bend" causes it to date the beginning of the productivity growth revival of the 1990s well before 1995, and it measures the productivity growth trend in 2002–03 at a relatively low 2.35 percent a year. As a compromise, I use hereafter the H-P filter with a parameter of 6,400, which produces a final mid-2003 value for trend productivity growth of 2.76 percent a year.

The bottom panel of figure 1 compares the H-P 6,400 trend with two versions of a Kalman trend based on a time-varying coefficient as in equation 5 above. As the additional explanatory variable set (X_t), I use the current and four leading values of the change in the output gap, which in turn is defined as the log ratio of real GDP to its TTB-based trend. I use leads rather than lags because a consistent feature of the data is that productivity leads output, which in turn leads aggregate hours of work; table 5 below regresses changes in hours on *lagged* changes in the output gap and regresses productivity on *leading* changes in the output gap.

The two Kalman trends in the bottom panel of figure 1 differ only in that one includes and the other excludes the output gap variable, thus allowing visual inspection of the impact of including the output gap as a piece of outside information that is unavailable to the H-P filter. The

12. Figure 1 extrapolates the 1994–2001 TTB trend to 2003:2, thus ignoring any information on rapid productivity growth actually achieved between 2001:2 and 2003:2.

Kalman version that omits the output gap is remarkably similar to the H-P 6,400 trend except in the last few quarters: it has a terminal value of 3.14 percent a year, considerably above the terminal H-P value of 2.74 percent a year. Both the no-gap Kalman trend and the H-P trend exhibit virtually identical plunges in trend productivity growth in the 1978–83 period, corresponding to a period of negative growth in the output gap, whereas the with-gap Kalman trend has a much smoother profile over the entire 1978–88 decade. Similarly, the no-gap Kalman trend exhibits a hump in the 1963–68 period, as does the H-P trend, whereas the Kalman trend using the output gap variable displays a smoother pattern, registering a slower trend in 1962–68 when the output gap was rising and a faster trend in 1968–76 when the output gap was declining.[13]

In developing the Kalman trends shown in the bottom panel of figure 1, I considered numerous values of the smoothing parameter and based the choice on two criteria.[14] First, the Kalman trend that omits the output gap variable behaves with uncanny similarity to the HP 6,400 trend, so that these two trends are consistent in their notion of smoothness. Second, the same Kalman smoothness parameter provides an optimistic interpretation of the 2002–03 surge in actual productivity growth and thus leans against the natural tendency to view the 2002–03 behavior skeptically, as an ephemeral event. The terminal value for the Kalman trend (using output) is 3.38 percent a year, and an alternative parameter that imposes additional smoothness yields a terminal value of 2.96 percent a year.[15]

Although the Kalman trend that includes the output gap variable might seem to be the best choice, it suffers from an important defect. Since the time-varying trend is estimated in an equation explaining the *growth rate* of productivity, the resulting growth trend provides no information on the *level* of the productivity trend or on the ratio of actual to trend at any par-

13. For the full period through 2002:2, the Kalman filter that uses the output gap is based on the current and four leading values. The trend is extrapolated through 2003:2 using an alternative version based on the current and four *lagged* values.

14. Estimation was performed using the RATS DLM procedure, using a routine that was developed for this purpose with the help of Tom Doan, author of RATS. The "sv" parameter that controls the variance of the "v" error term in equation 6 was set alternatively at values of 1, 2, 4, 8, 16, and 32. The plots in the bottom panel of figure 1 and the Kalman trends used in the rest of the paper are based on setting sv equal to 16.

15. The Kalman trends in figure 1 are based on a smoothness parameter sv = 16, and the alternative result cited in the text, which imposes greater smoothness, is estimated with sv = 32.

ticular time. Thus a central question—how far from the actual value was trend productivity in mid-2003?—cannot be answered with this technique. The H-P filter, in contrast, is estimated on the *level* of productivity and thus generates both the level and the growth rate of trend productivity automatically, and, as the bottom panel of figure 1 shows, it estimates a substantially lower trend (and a higher ratio of actual to trend) in mid-2003 than does the Kalman filter.[16]

In the end I have chosen not to rely exclusively on the Kalman filter version with the output gap variable but rather to create a trend that is an average of this Kalman trend and the H-P 6,400 trend shown in the same figure. To determine the *level* of the trend, I have imposed the constraint that the mean log deviation of actual productivity from trend over the period 1954:4 to 2003:2 is exactly zero; this constraint yields three series for the deviation of actual productivity from trend, shown in figure 2: one using the with-output Kalman series in figure 1, one based on the H-P 6,400 trend, and the average of the two. The time paths of the two alternative ratios of actual to trend are very similar, and because the H-P trend tends to respond to cyclical movements in the actual value more than the Kalman trend, the actual-to-trend ratio for the H-P series in figure 2 tends to lie "inside" the Kalman series (that is, closer to zero), as is evident in 1963–69, 1979–83, and 1993–2001.
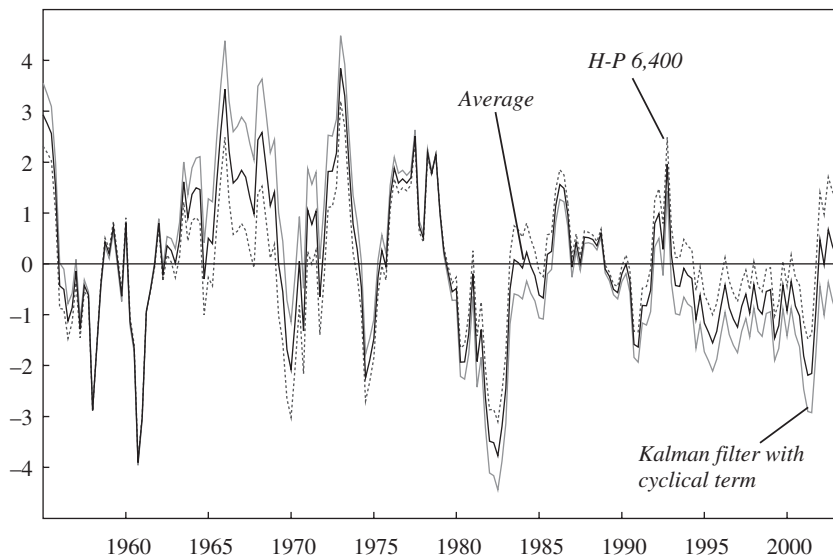
Perhaps the most surprising conclusion from figure 2 is that both trend methods imply that the level of actual productivity was below trend throughout the entire period 1993–2001; the H-P ratio of actual to trend rises substantially above zero only in 2002–03, and the Kalman ratio is negative over the entire period between early 1993 and mid-2003. These results refute my earlier position that actual productivity was above trend in the late 1990s;[17] that position was defensible based on data at the time, but the soaring growth of actual productivity in 2002–03 now pulls the estimated trend lines well above actual values for the late 1990s.

16. Considerable effort was devoted to estimating a level model for the Kalman filter in which the log-level of productivity was regressed on a time trend with a time-varying coefficient, on a constant, and on various combinations of the level and change of the output gap. The estimated trend growth rate of productivity had much too little variance, fluctuating between only 1.75 percent and 2.25 percent over the full forty-eight-year sample period, even at very low smoothness parameters. I regard those results as implausible, but further exploration of this issue would require building a more complex time-series model, which is beyond the scope of this paper.

17. Gordon (2000).

Figure 2. Deviation of Actual Productivity from Alternative Specifications of Trend, 1955–2003[a]

Percent (log ratio)

Source: Author's calculations.
a. Trend estimation methods are those depicted in the bottom panel of figure 1.

For symmetry, the other components of the output identity are detrended by the average of the Kalman and H-P growth rate trends, using the same method and parameters. To save space, I do not plot the growth trends of these other components; the trends rise and fall smoothly, reflecting movements of growth rates over long periods that are summarized by the TTB trends in table 1. The single exception to using the average of the Kalman and H-P trends involves the employment rate. The trends should reflect what is meant by changes in potential output, or what the economy can produce when operating at the NAIRU. The NAIRU changes slowly in response to structural changes in labor market behavior and does not suddenly jump up or down in response to booms or slumps. The counterpart of the NAIRU in this analysis is the trend employment rate, which is simply one minus the NAIRU. This should evolve independently of the business cycle, and no smoothing parameter for the Kalman or H-P filter was found to achieve the desired degree of stability; in particular, the H-P 6,400 trend employment rate falls below 93 percent from

1980 to 1988 and to almost 92 percent in 1983, implying a NAIRU of 8 percent.

The desired stable behavior of the trend employment rate is, however, achieved when the TTB method is used, yielding the changes in the trend employment rate illustrated in table 1.[18] The TTB trend is almost flat (table A1 in appendix A lists the unemployment rates in each benchmark quarter): it exhibits a slight decline between 1955 and 1978, the counterpart of the gradual upward creep of the NAIRU in that era, then stabilizes at about 94 percent through 1987 (the counterpart of a 6.0 percent NAIRU), and then gradually increases to 95 percent by 2001. This last is the counterpart of the decline of the NAIRU in the 1990s from 6 percent to 5 percent, which was evident in the early analyses of the time-varying NAIRU and has been validated since then.[19]

Summary of Trend Choice and Implications for Real GDP

This discussion of the choice of methodology for trend estimation concludes that the average of a time-varying coefficient estimated with the Kalman filter and that estimated with the H-P filter with a parameter of 6,400 (instead of the conventional 1,600) provides nicely smoothed trend rates of change for six of the seven components of the right-hand side of the output identity. The level of the trend for each variable is established by imposing the constraint that the average ratio of actual to trend is equal to zero over the full period 1954:4 to 2003:2. The exception is the employment rate, where the TTB method is used to establish a trend series that has the desired resemblance to the NAIRU as estimated in recent studies of inflation behavior.

For two labor market variables, the employment rate and the LFPR, there is a special treatment of the trend in the two years after the final benchmark quarter (2001:2). In each case the trend is set equal to its value

---

18. The TTB method yields a trend employment rate series that is very close to the alternative of using a NAIRU series like that estimated by Eller and Gordon (2003). Since the two series are so close, simplicity dictates using the TTB method, already described here, rather than the Eller-Gordon series, which would require considerable additional explanation.

19. Early estimates of the time-varying NAIRU appeared in the same issue of the *Journal of Economic Perspectives* (Staiger, Stock, and Watson, 1997; Gordon, 1997). Updated estimates showing that the NAIRU leveled off at values slightly above 5 percent after 1999 are reported by Eller and Gordon (2003) and Staiger, Stock, and Watson (2001).

in that quarter; that is, its growth rate is zero. In each case this special treatment is justified to avoid treating the cyclical decline in these labor market variables as a trend rather than a cyclical phenomenon. Special treatment is not required in earlier episodes.

The trend in real GDP is simply the sum of the log trends in the seven component trends for the variables on the right-hand side of equation 3. The top panel of figure 3 shows the eight-quarter growth rate of actual real GDP and its trend, revealing that the real GDP trend avoids both the sharp step jumps of the TTB trend for real GDP exhibited in table 1 and the excessive waviness associated with H-P trends. The bottom panel of figure 3 exhibits the log ratio of actual to trend real GDP, which clearly distinguishes the major output losses of 1981–83 from the intermediate episodes of 1957–58, 1960–61, and 1975–76 and the milder episodes of 1970–71, 1990–91, and 2001–03.
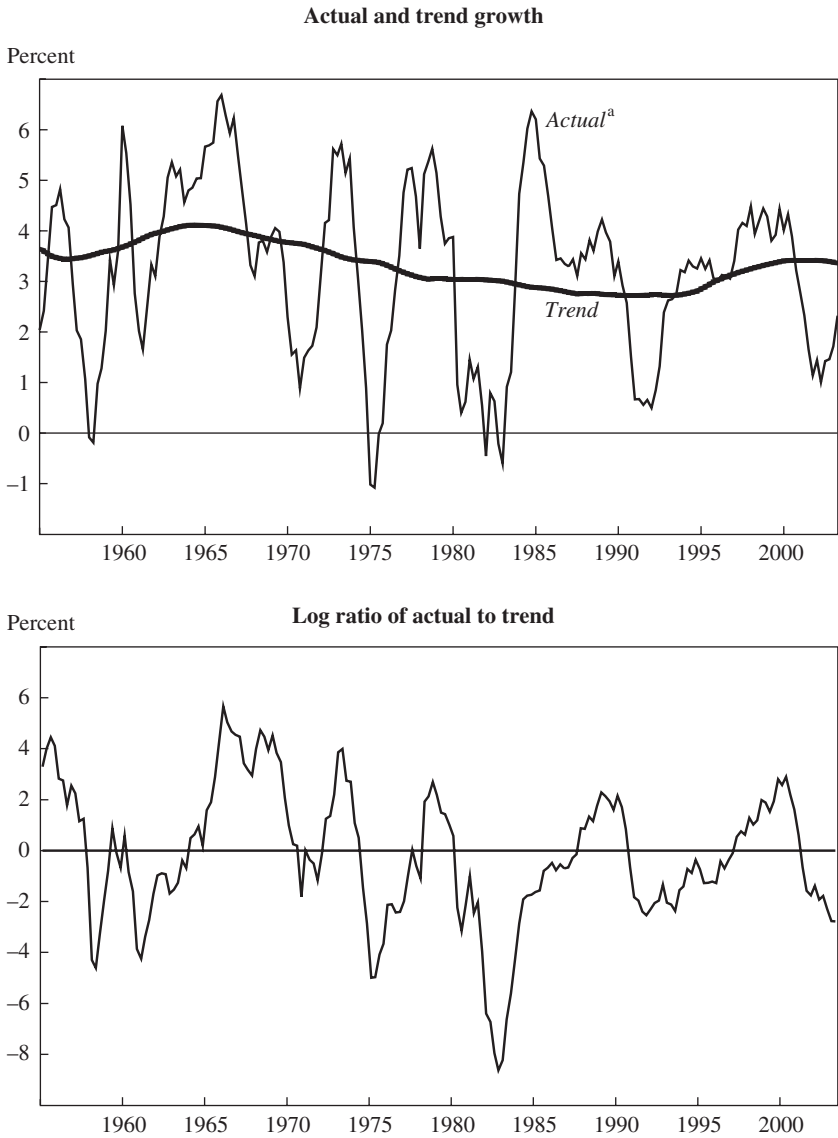
A New Decomposition Based on Okun's Law

Before turning to a regression analysis that takes account of lagged reactions, I employ a crude but revealing technique to determine the contributions of the components of the output identity to cyclical fluctuations in output. For each of the periods between benchmark quarters, I determine the quarter when actual real GDP reached its highest log ratio to trend GDP and the quarter when that ratio was the lowest. Table 3 lists these peak and trough quarters along with the log ratios of actual to trend in those quarters for real GDP and the seven right-hand-side components of the output identity.

An Examination of Ratios of Actual to Trend

Does the actual-to-trend ratio for the employment rate equal roughly one-third of the actual-to-trend ratio for real GDP, as Okun's original analysis would suggest? With two exceptions, the employment response in peaks lies in a narrow range between 30 and 37 percent of the actual-to-trend ratio for real GDP (fourth column of table 3 divided by the first column), not far from Okun's initial value. But in troughs the response of unemployment is much greater, ranging from 45 to 59 percent, with the single exception of 1970:4. Taking all episodes, the average response is

Figure 3. Actual and Trend Growth in Real GDP and Their Log Ratio, 1955–2003

Actual and trend growth

Percent

Log ratio of actual to trend

Percent

Source: Author's calculations.
a. Eight-quarter change.

Table 3. Peak and Trough Values of Actual-to-Trend Log Ratios of Components of Real GDP Growth, 1955–2003

Percent

Quarter	Real GDP	Output per hour[a]	Hours per employee[a]	Employment rate	LFPR	Working-age population	Mix effect[b]	Employment ratio[b]
Peak quarters								
1955:3	4.45	2.56	0.76	1.33	0.57	–0.24	–0.72	0.18
1966:1	5.67	3.44	1.74	1.78	–0.60	–0.07	–0.69	0.08
1973:2	3.99	3.30	0.58	0.87	–0.14	0.14	–1.25	0.49
1978:4	2.69	2.15	0.33	0.15	0.91	0.36	–1.41	0.21
1989:1	2.29	–0.10	0.52	0.76	0.75	0.19	0.09	0.08
2000:2	2.90	–0.35	0.75	1.08	0.58	0.09	–0.41	1.16
Trough quarters								
1958:2	–4.60	–1.63	–1.26	–2.08	0.59	–0.22	1.31	–1.31
1970:4	–1.82	–1.32	–0.96	–0.17	0.30	–0.29	0.80	–0.18
1975:1	–5.00	–1.47	–1.37	–2.59	–0.39	0.22	0.46	0.13
1982:4	–8.63	–3.20	–1.34	–5.08	–0.05	0.38	1.73	–1.07
1991:4	–2.54	–0.52	–1.07	–1.41	–0.29	–0.06	0.86	–0.05
2003:2	–2.78	0.94	–0.47	–1.37	–0.44	0.04	–0.52	–0.96
Average at peaks	3.67	1.84	0.78	0.99	0.34	0.08	–0.73	0.37
Average at troughs	–4.23	–1.20	–1.08	–2.12	–0.05	0.01	0.77	–0.57
Average, peaks and troughs (absolute value)	3.95	1.52	0.93	1.56	0.20	0.03	–0.75	0.47
Share of total change in real GDP	100	38.4	23.5	39.4	4.9	0.8	–19.1	11.9

27 percent for peaks and 50 percent for troughs, or 39 percent when peaks and troughs are averaged together (bottom line of table 3, fourth column).

Which other components of the output identity are most important in accounting for output fluctuations? The bottom line of table 3 shows that productivity makes almost the same contribution as the employment rate, with a response of 38 percent; hours per employee contributes 24 percent, and the LFPR only 5 percent. The mix and payroll-household employment ratio effects together contribute –7 percent, for a total response (net of rounding error) of 100 percent. These averages disguise varying contributions in particular episodes. In two peak and two trough episodes, the LFPR contribution has the wrong sign, helping to account for its small response in the overall average. In the last two peak episodes, the productivity contribution is negative, but in the final trough episode its contribution is strongly positive. Table 3 thus suggests that sometime after the mid-1980s the role of productivity in output cycles changed significantly.

A Dynamic Specification for the Components of the Output Identity

My primary interest in developing a dynamic specification suitable for regression analysis is to provide the best possible representation of average cyclical responses of the components of the output identity across forty-eight years of history. This will be the perspective taken subsequently to assess unusual aspects of the expansion of the late 1990s, the period of recession and early recovery after 2000, and, by way of comparison, the early recovery period of 1991–92.

CHANGES IN DEVIATIONS FROM TREND AND AN ERROR CORRECTION TERM. The point of departure for the dynamic specification here is work by Christopher Sims and my own earlier work on cyclical productivity issues.[20] The dependent variable is each of the components of the right-hand side of the output identity (equation 2), which is transformed into log values as in equation 3. Each dependent variable is expressed as the first difference of the logarithm of the variable, say, x, and the logarithm of its trend x^*; in the notation introduced above, this is $\Delta x_t'$. This variable is regressed on a series of lags of the dependent variable and on the first differences of deviations of the logarithm of real GDP from its trend ($\Delta q_t'$). The output deviation variable can in principle enter with leads, the

---

20. Sims (1974); Gordon (1979, 1993).

current value, and lags. The lags can be interpreted as reflecting adjustment costs and, for such components as the employment rate and the LFPR, delays in hiring and firing. The use of leads was introduced by Sims in his analysis of Granger causality between hours and output.[21] I provide a separate treatment of the productivity component of the output identity, specifying the productivity-to-output relation alternatively as a regression with hours lagging behind output,[22] or as productivity leading output.

Two additional variables are added to the traditional regression that relates the first differences of component deviations ($\Delta x_t'$) to first differences of output deviations ($\Delta q_t'$). The first is an error correction term. The concept of error correction has been linked to that of cointegration, which can be defined informally as the notion that a linear combination of two series—for example, the hours deviation and the output deviation—is stationary.[23] When two such variables are cointegrated, a regression consisting entirely of differenced data will be misspecified, whereas a regression consisting entirely of level data will omit important constraints. The solution is to estimate a regression of the first difference of one variable on the first difference of the other, adding to the regression an error correction variable consisting of the lagged log ratio of one variable to the other.[24] In my application of this technique, I impose stationarity on the error correction term by entering it as the lagged log ratio of actual to trend of the variable in question, whether it is productivity, the employment rate, or one of the other components of the identity. In summary, this specification explains the rate of change of a deviation from trend by the rates of change of the deviation from trend of the lagged dependent variable and of output, and the lagged *level* of the deviation of the dependent variable from its own trend.

THE END-OF-EXPANSION EFFECT. In my 1979 paper, verified and extended in 1993, I identified a tendency in the late stages of a business expansion for labor input to grow more rapidly than can be explained by

---

21. Sims (1974).

22. As in Gordon (1993).

23. For a formal definition of stationarity and cointegration, see Engle and Granger (1987, pp. 252–53).

24. A complete taxonomy of the possible forms of dynamic specification in a bivariate model is presented in Hendry, Pagan, and Sargan (1984, pp. 1040–49).

output changes.[25] I dubbed this tendency toward overhiring the "end-of-expansion" (EOE) effect and argued that it was balanced by a tendency to underhire in the first two years or so after the end of the expansion. The dynamics of the EOE effect reflect overoptimistic forecasts and overhiring as the economy's growth slows and finally halts at the end of an expansion, causing a squeeze on profits, followed by cost cutting, layoffs, and a reluctance to hire workers in the recession and first stages of the recovery.

Although originally developed to explain cyclical productivity behavior, the same phenomenon should show up in equations for some or all of the labor market variables. If overhiring holds productivity down at the end of expansions, that same overhiring should be evident in some combination of the employment rate, the LFPR, hours per employee, and perhaps even the mix and payroll-household employment ratio terms. It is particularly interesting to reassess the EOE effect in light of the apparent change in cyclical behavior of productivity since the mid-1980s previously discussed in connection with table 3.

The EOE effect is introduced into the dynamic specification through a set of seven dummy variables, corresponding to seven EOE episodes since 1955. These are not conventional dummies that take values of zero or 1; rather they are in the form $1/M$ and $1/N$, where M is the length in quarters of the initial interval of excessive labor input growth, and N the length of the subsequent correction. By forcing the sum of the coefficients on each variable to equal zero, any overhiring in the initial phase is subsequently corrected. Overhiring that is *not* balanced by subsequent underhiring will result in a small and insignificant coefficient on the EOE dummy and will come out either in the equation's residual or in the coefficients on other variables.

In my 1993 paper I determined the dating of the EOE dummies by referring to the distinction between the business cycle as determined by the National Bureau of Economic Research (NBER) and the growth cycle. According to the NBER definition, an expansion ends when real output reaches its absolute peak. This can be distinguished from the earlier peak of the growth cycle, which occurs when output reaches its highest level relative to trend or potential output. My 1993 paper set the first M

25. Gordon (1979, 1993).

quarters as the period between the peak in the growth cycle and the peak of the NBER cycle. The timing and duration, *N,* of the subsequent correction period is determined by examining residuals in equations that omit the dummies entirely. The amplitude of the EOE effect is allowed to differ across business cycles by allowing the dummy variable for each cycle to have its own coefficient; subsequently I test whether these coefficients can be pooled into a single coefficient.[26]

Combining these explanatory variables, the basic equation to be estimated for the components of the output identity is

$$(7) \qquad \Delta x'_t = \sum_{i=1}^{4} \alpha_i \Delta x'_{t-i} + \sum_{j=0}^{4} \beta_j \Delta q'_{t-j} + \phi x'_{t-1} + \sum_{k=1}^{7} \gamma_k D_k + \varepsilon_t,$$

where $D_k = 0$ in all quarters except the EOE and subsequent correction period, which are as listed in table A2 in appendix A. Here the α_i are the coefficients on the lagged dependent variable; the β_j are the current and lagged coefficients on the change in the real GDP deviation from trend; ϕ is the coefficient on the error correction term; and the γ_k are the coefficients on the EOE dummies. The γ_k coefficients are interpreted according to which dependent variable is being explained; labor market variables such as hours of labor input or the employment rate would be expected to have a positive EOE coefficient, whereas a regression with productivity as the dependent variable would be expected to have a negative coefficient.

Estimation for Six Components of the Output Identity

The results from estimating equation 7 are presented in tables 4 and 5. First, table 4 presents results for five of the identity components, that is, all of them except for productivity and population. Regressions for population are omitted because there is no reason to expect the working-age population to respond cyclically to changes in output, and indeed preliminary testing indicated no response, with adjusted R^2s in the vicinity of 0.01. Regressions for productivity are treated in table 5 and are given special attention, in view of the topical interest in the performance of produc-

26. Gordon (1993, p. 291, notes 33 and 34) discusses several arbitrary choices that were made in carrying out this definition of the *M* quarters. This paper takes the definition of the EOE dummies from Gordon (1993), with a few minor changes of a quarter more or less. The *M* quarters for the late 1980s are identical to the "early" EOE dummy discussed in Gordon (1993, p. 300).

Table 4. Regressions Explaining Cyclical Deviations from Trend in Output Identity Components[a]

Independent variable	*Hours per employee*[b]	*Employ-ment rate*	*LFPR*	*Mix effect*[c]	*Employ-ment ratio*[c]
	Dependent variable				
Lagged dependent variable[d]	–0.17	–0.23	0.18	–0.73**	–0.09
Output deviation from trend[d]	0.29**	0.50**	–0.00	–0.32**	0.16**
Error correction term[e]	–0.85**	–0.11*	–0.92**	–0.06	–0.34*
End-of-expansion (EOE) dummy variables[f]					
1954–58	0.51	1.18**	0.11	0.33	1.30*
1059–62	0.42	0.19	0.85	0.64	–0.24
1968–71	0.74	0.97**	0.77	0.96	0.84
1973–76	0.73	1.30**	0.67	0.30	0.40
1978–83	–0.23	0.58	0.90	1.11*	0.51
1988–92	0.90	0.55	0.88	0.74	0.65
2000–03	–0.07	0.56	0.62	0.72	1.19*
Summary statistics:					
Adjusted R^2	0.40	0.73	0.14	0.42	0.16
Standard error of estimate	1.02	0.79	1.02	1.21	1.28
Sum of squared residuals	180	108	180	255	287
Addendum: all EOEs are constrained to be equal					
EOE	0.39**	0.75**	0.67**	0.67**	0.73**
Summary statistics:					
Standard error of estimate	1.01	0.79	1.00	1.20	1.28
Sum of squared residuals	185	112	182	257	294

Source: Author's regressions using equation 7 in the text and data from sources and methods described in appendix A.

a. * indicates coefficient or sum of coefficients is statistically significant at the 5 percent level, ** indicates significance at the 1 percent level.

b. Data are for the nonfarm private business sector.

c. Defined as in table 1.

d. Four lags are used.

e. Lagged log ratio of actual to trend of the dependent variable.

f. Values of the variables are in the form 1/*M* or 1/*N*, where *M* is the length in quarters of the initial interval of excessive labor input growth, and *N* the length of the subsequent correction. See appendix A.

tivity during 2000–03 and the need for a forecast for productivity growth
if one is to make projections of potential GDP growth over the next two
decades.

The five dependent variables are each expressed as in equation 7 as the
first difference of a log ratio of the actual value to the trend value. The
coefficients are presented in rows corresponding to their order in equa-
tion 7, and the bottom three lines of the table provide alternative estimates
(and summary statistics) of the γ_k coefficients for the EOE effect that
impose the constraint that all of these seven coefficients are equal.

Let us start with the employment rate results in the second column of
table 4. The sum of coefficients on the output deviation is 0.50, substan-
tially above the 0.39-percentage-point effect reported in the bottom line
of table 3. The error correction term has the expected negative sign, indi-
cating that a high value of the lagged ratio of the employment rate to its
trend tends to push down the growth rate of the employment rate relative
to its trend. All of the EOE coefficients have the expected positive sign,
and three are highly significant, indicating that the employment rate was
pushed up by overhiring in the late stages of expansion in 1955–57,
1968–69, and 1973–74 but was then pushed down by underhiring.

Scanning across the other four columns reveals a strong positive
response of hours per employee and of the payroll-to-household employ-
ment ratio, and a strong negative response of the mix effect to cyclical
output movements, but no significant response of the LFPR. All the error
correction terms in these columns are significant, with the expected nega-
tive sign, except for the mix effect. Only scattered individual EOE coeffi-
cients are significant. However, when the EOE coefficients are forced to
be equal, as at the bottom of the table, they are all significant and have the
expected positive sign. I return to the implications of the EOE coefficients
in table 4 after looking at the productivity results in table 5.

Dynamic Equations for Aggregate Hours and Productivity

In my previous papers,[27] empirical equations did not explain labor pro-
ductivity changes themselves but rather took the quarterly change in the
log ratio of actual to trend aggregate hours of work as the dependent vari-
able. Recall that, in the notation of equation 2 above, output per hour is
Q^B/A^B. The denominator is aggregate hours of work, that is, NFPB employ-

---

27. Gordon (1979, 1993).

ment times NFPB hours per employee. In this paper I estimate similar equations explaining aggregate hours and compare the results with those in my 1993 paper but also estimate symmetric equations that explain the quarterly change in the log ratio of actual to trend output per hour.

A familiar aspect of productivity dynamics is that A^B responds with a lag to cyclical movements in Q^B, and this lagged adjustment of A^B implies that productivity, the ratio Q^B/A^B, *leads* output movements. Although it is no problem to run a regression for productivity change with the specification of equation 7 in which leads on the output deviation term replace lags, this has the practical disadvantage that it prevents one from estimating the equation or providing residuals that cover a period of great interest, namely, the four quarters ending in 2003:2. In order to estimate productivity equations that can provide predicted values and residuals for these final four quarters, one can omit the leading values of the output deviation variable and include only the current value, or the current value along with several lags. As in table 4, the explanatory variables include four lags on the dependent variable, the current and four leads or lags on the change in the output deviation, and the error correction term defined as the lagged log ratio of actual to trend productivity. Since one would expect above-trend productivity to stimulate hiring, the error correction term should have a positive coefficient when aggregate hours is the dependent variable, and a negative coefficient when productivity itself is the dependent variable.[28]

Column 5-1 of table 5 reproduces results from my 1993 paper for a sample period ending in 1992:4.[29] These are then compared in column 5-2 with the results for the same sample period and specification using currently available data. The coefficients in the first two columns should differ only as a result of data revisions in the intervening decade. These data revisions have indeed introduced several changes in the results. They have made the data more noisy, as indicated by the deterioration in goodness of fit. The sum of the coefficients on the lagged dependent variable and on the output deviation becomes larger in absolute value, and the error correction term, although also larger, is still insignificant. The EOE

---

28. Since the dependent variable in the aggregate hours equation is the change in the log output ratio of actual to trend, a trend for aggregate hours is needed. Appendix A describes how it is developed.

29. Gordon (1993, table 5).

Table 5. Regressions Explaining Cyclical Deviations from Trend in Aggregate Hours Worked and Productivity[a]

	Dependent variable					
	Aggregate hours				Output per hour	
Independent variable	5-1[b]	5-2[c]	5-3[d]	5-4[d]	5-5[d]	5-6[d]
Lagged dependent variable[e]	–0.29*	–0.44**	–0.32*	0.06	–0.50**	–0.19
Output deviation from trend[e]	0.91**	1.13**	1.07**	0.81**	0.50**	0.71**
Error correction term[f]	0.08	0.28	0.27*	0.45**	–0.36**	–0.81**
End-of-expansion (EOE) dummy variables[g]						
1954–58	2.25**	3.28**	2.99**		–3.75**	
1959–62	1.95**	2.03*	1.92*		–2.47**	
1968–71	2.71**	2.82**	2.53**		–4.03**	
1973–76	3.24**	3.26**	2.92**		–3.62**	
1978–83	2.65**	0.71	0.53		–2.61**	
1988–92	3.01**	2.65**	2.34**		–3.39**	
2000–03			1.47*		–1.76*	
Summary statistics:						
Adjusted R^2	0.83	0.79	0.76	0.72	0.73	0.66
Standard error of estimate	1.50	1.66	1.63	1.79	1.77	2.00
Sum of squared residuals	309	366	464	581	533	706
Addendum: all EOEs are constrained to be equal						
EOE	2.57**	2.40**	2.02**		–3.00**	
Summary statistics:						
Standard error of estimate	1.49	1.68	1.65		1.77	
Sum of squared residuals		389	488		550	

Source: Author's regressions using equation 7 in the text and data from sources and methods described in appendix A.

a. * indicates coefficient or sum of coefficients is statistically significant at the 5 percent level, ** indicates significance at the 1 percent level. Sample period for all regressions begins in 1954:4.

b. Results reported in Gordon (1993) using then-current data through 1992:4.

c. Same specification as in Gordon (1993) but using currently available (revised) data through 1992:4.

d. Same specification but using currently available data through 2003:2.

e. Four lags of the dependent variable are used, except in columns 5-5 and 5-6, where four leads are used instead.

f. Lagged log ratio of actual to trend of the dependent variable.

g. Values of the variables are in the form $1/M$ or $1/N$, where M is the length in quarters of the initial interval of excessive labor input growth, and N the length of the subsequent correction. See appendix A.

coefficient for the 1978–83 period declines sharply and loses statistical significance.

Column 5-3 extends the sample period by forty-two quarters, to 2003:2. The results change remarkably little: the coefficients are about the same, the standard error actually declines somewhat, and the EOE coefficient for 2000–03 provides a new example of the continued relevance of the EOE hypothesis, first developed almost twenty-five years ago. Column 5-4 uses an identical specification and sample period except for the omission of the EOE dummies, which raises the standard error of the estimate from 1.63 to 1.79 and the sum of squared residuals by 25 percent.

Column 5-5 covers the same time period as columns 5-3 and 5-4 but switches the dependent (and the lagged dependent) variable from the change in the hours deviation from trend to the change in the productivity deviation from trend, still within the specification of equation 7. The four lags on the output deviation are replaced with four leads. The productivity version in column 5-5 fits somewhat less well than the hours version, which is not surprising since productivity changes are more variable than changes in hours. Comparing column 5-5 with column 5-3, one observes that the use of labor productivity instead of hours as the dependent variable raises the statistical significance of all variables and increases the absolute value and the significance of all seven EOE dummy variables. As expected, the coefficients on the error correction term and the EOE dummies have opposite signs in columns 5-3 and 5-5.[30]

How should the coefficients on the output deviation be interpreted in the hours and productivity versions? The short-run response of hours to output is 1.07 in column 5-3, and this translates into a long-run response of 0.81 when divided through by 1 minus the sum of coefficients on the lagged dependent variable. This, in turn, gives a response of productivity change to output change of 1 – 0.81, or 0.19.[31] The same arithmetic implies a long-term response in column 5-5 of 0.33, and this larger response presumably reflects to some extent the simultaneity of the output and productivity variables. That is, running a regression of productivity change on current output change creates an upward bias on the output coefficient in column 5-5.

30. The leads are highly significant. The $F(4,181)$ ratio for the inclusion of the four leads is 6.03, compared with a 5 percent critical value of 2.42.

31. The long-run responses of 0.19 in column 5-3 and 0.33 in column 5-5 bracket the long-run response of 0.28 reported in Gordon (1993, p. 294).

Table 6. Short- and Long-Run Responses to Output Changes of Output Identity Components, 1954:4–2003:2

Component	From table 3	Estimate A[a]		Estimate B[b]	
		Short run	*Long run*[c]	*Short run*	*Long run*[c]
Output per hour[d]	0.38	−0.07	0.19	0.50	0.33
Employment rate	0.39	0.50	0.41	0.50	0.41
Labor force participation rate	0.05	0.00	0.00	0.00	0.00
Working-age population	0.01	. . .	. . .	. . .	. . .
Hours per employee[d]	0.23	0.29	0.25	0.29	0.25
Mix effect	−0.19	−0.32	−0.18	−0.32	−0.18
Payroll-household employment ratio	0.12	0.16	0.15	0.16	0.15
Sum of coefficients	0.99	0.56	0.82	1.13	0.96

Sources: Tables 3, 4, and 5 and author's calculations.

a. Estimate for the short-run response for output per hour is calculated from the coefficient on output deviation from trend from the aggregate hours equation (equation 5-3) in table 5; other short-run responses from table 4.

b. Estimate uses the coefficient on output deviation from trend from the output per hour equation (equation 5-5) in table 5 for the short-run response for output per hour; other short-run responses from table 4. Sample period ends in 2002:2.

c. Sum of coefficients on the output term divided by 1 minus the sum of coefficients on the lagged dependent variable.

d. For the nonfarm private business sector.

The bottom three lines of table 5 again report results of a regression that forces the separate EOE coefficients to be identical across the six or seven EOE episodes. In each case one cannot reject the hypothesis that all the EOE coefficients are identical, and the case for equal coefficients is strongest in column 5-5.[32] A nicely symmetrical result is that when the EOE coefficients are forced to be equal for the components of the output identity in table 4, the sum of these EOE effects across the five columns of the table is 3.21, which is very close to, and with the opposite sign from, the combined EOE effect on productivity of −3.00 at the bottom of column 5-5 in table 5.

How do the long-run responses of the components of the output identity to cyclical movements in output estimated in tables 4 and 5 compare with the crude peak and trough ratios developed in table 3? These results are brought together in table 6, which repeats the response coefficients from table 3 and both the short-run and the long-run responses derived from tables 4 and 5, shown alternatively for the versions of the equation

---

32. The $F(5,138)$ ratio in column 5-2 is 2.06, compared with a 5 percent significance level of 2.27. The $F(6,180)$ ratio in column 5-3 is 1.55, compared with a 5 percent significance level of 2.14. And the $F(6,180)$ ratio in column 5-5 is 0.45, compared with the same 5 percent significance level of 2.14.

using hours as the dependent variable (estimate A) and that using productivity as the dependent variable (estimate B). The response coefficient of the employment rate is quite stable across the columns of table 6, varying only from 0.39 to 0.50. Similarly, the coefficient on hours per employee is quite stable in a range of 0.23 to 0.29. Although the productivity responses appear to be quite inconsistent across columns, the long-run response of 0.33 in the last column of table 6 is not far from the response of 0.38 listed in the first column. The columns agree that the LFPR effect is unimportant, and the mix and payroll-household employment ratio effects taken together seem to suggest a long-run negative response of about –0.03.

Productivity Growth in 2000–03: Trend, Cycle, and Residual
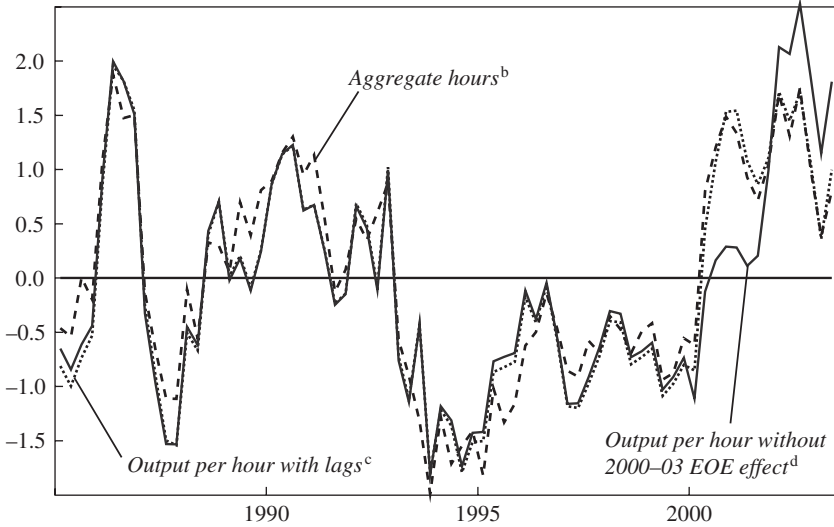
The above analysis provides the tools to determine how much of the rapid productivity growth since 2000 was predicted by the regression equation and how much emerges in the unexplained residual. I examine the results in three steps. First, figure 4 presents a plot of productivity growth residuals. Then, table 8 below compares the behavior of actual and predicted productivity growth in the initial 2002–03 recovery phase with that in the early stages of six previous business cycles. Finally, table 9 reports actual and predicted changes in all the components of the output identity over selected intervals, including 2000–03, showing to what extent productivity residuals are balanced by residuals in the opposite direction in other components of the output identity.

Figure 4 plots four-quarter averages of residuals from three different versions of the regression equation as specified in table 5. The line labeled "Aggregate hours" plots the residuals for the equation explaining the change in the hours deviation in column 5-3 of table 5, which includes the current value of the change in the output deviation and four lags. The figure reverses the signs of the residuals in the hours equation, because a negative residual for the change in detrended hours implies a positive residual in the change for detrended productivity.

The line labeled "Output per hour with lags" plots the residuals for the equation explaining the change in the productivity deviation in table 5, column 5-5, except that for this purpose I omit the lead terms on the output deviation in order to be able to compute residuals through to the final

Figure 4. Residuals from Productivity Growth Equations, 1985–2003[a]

Percent a year

Source: Author's calculations.
a. All residuals are four-quarter averages.
b. Residuals from equation reported in column 5-3 of table 5 (hours is the dependent variable).
c. Residuals from equation reported in column 5-5 of table 5; leads are replaced by lags to produce values through 2003:2.
d. Same specification as output per hour with lags, except that the final EOE dummy variable for 2003–03 is omitted.

quarter 2003:2, and instead substitute four lags.[33] These two lines are almost identical, as they should be, and indicate a tendency of the specification to underpredict actual productivity growth during 1988–93, to overpredict it during 1994–99, and to underpredict it substantially during 2000–03. It is interesting to compare the residuals of the second equation in the 2000–03 episode with those in the previous recession and recovery in 1990–93. To do this I designate the recession quarters as beginning with the output peak and extending three quarters beyond, and the recovery quarters as the first seven quarters after the trough in real GDP. Table 7 reports the average residual during each interval for both the pro-

    33. When four leads are included as in the version shown in table 5, column 5-5, the sample period of the regression must stop in 2002:2, one year before the end of the data. Substituting four lags for four leads raises the sum of the squared residuals (SSR) in column 5-5 from 533 to 576, and including only the current value with no leads or lags raises the SSR further to 604.

Table 7. Average Residuals of Productivity and Hours Equations in Recent Cyclical Episodes

Percent

Episode	Productivity equation[a]	Hours equation[b]
Recessions		
1990:2–1991:1	0.92	1.13
2000:4–2001:3	1.25	0.72
Recoveries		
1991:2–1992:4	0.29	0.39
2001:4–2003:2	1.20	1.14

Source: Author's calculations.

a. Residuals from equation specified as in column 5-5 of table 5 except that data are through 2003:2 and four lags of the output deviation variable are substituted for four leads.

b. Residuals from equation specified as in equation 5-3 of table 5 but with data through 2003:2.

ductivity equation (with lags on the output deviation rather than leads) and the hours equation.

The residuals in the productivity and hours equations are similar to each other in each of the four episodes. Both recessions share in common that productivity growth was stronger than the equation can explain—another sign of the tendency, already noted in table 3, for the behavior of productivity to become less procyclical in the last two cyclical episodes. The big difference in these two business cycles is that the residual in the first seven quarters of the 2001–03 expansion is much larger than that in the first seven quarters of the 1991–92 expansion.

The line labeled "Output per hour without 2000–03 EOE effect" in figure 4 displays an alternative version of the productivity equation that is otherwise identical to the first except that the final EOE dummy variable covering the 2000–03 period is omitted from the equation (the six EOE dummies for the earlier periods are retained). The difference between the two lines illustrates the size and timing of the estimated EOE effect in the 2000–03 episode. The full equation including the EOE effect predicts substantially less productivity growth in 2000 and early 2001 and substantially more during 2002–03 than the equation that omits the EOE effect, and so it has larger positive residuals in 2000–01 and smaller positive residuals in 2002–03. Intuitively, the EOE coefficient "sees" the robust productivity growth of 2002–03, interprets that as a typical EOE reversal impact, and—because the sum of the EOE initial and reversal effects is constrained to be zero—it is forced to estimate a negative EOE impact on productivity growth in the initial period 2000–01 even though

no such negative effect actually occurred. Taking the estimated EOE effect at face value, the unexplained component of productivity growth began to take off in mid-2000, about five quarters before the take-off became visible in the raw data. Recall that all these residuals are calculated for the deviation of productivity growth from a trend that is estimated to grow at a *much* faster rate than the consensus among academic economists only a few months ago: trend growth as estimated by the average of the H-P and Kalman filter results is 2.84 percent a year in 2000, 3.02 in 2001, 3.08 in 2002, and 3.05 in the first half of 2003.

The Early-Recovery Productivity Growth "Bubble"

Productivity growth was extremely rapid in the initial quarters after the trough of real GDP in 2001:3. But this is nothing new. A similar period of rapid productivity growth has been observed in the first few quarters of almost every postwar recovery, and in every case it has been followed by a return of productivity growth back to trend growth or even below trend growth in the subsequent phase of the expansion. This "early-recovery productivity growth bubble" is consistent with the EOE effect, for in the first few quarters of the recovery profits are still squeezed, and business firms are aggressively attempting to cut costs by reducing labor input.

Table 8 summarizes this little-noted cyclical phenomenon. The top panel reports, in the first column, the average growth rate of actual productivity relative to trend in the first four recovery quarters after each trough quarter for real GDP in seven cyclical recoveries going back to 1958. The bottom panel reports the same measure for the eight quarters following the first four quarters of the recovery. At the bottom of each panel is the average over the seven cyclical episodes. The remaining columns report averages over each interval of the predicted value from the productivity regression in table 5, column 5-5; the statistical contribution of the EOE dummy variables to that prediction; the prediction minus the EOE contribution; and finally the statistical residual.

The bottom line in each panel of table 8 shows the averages over the seven episodes in that panel and reveals a striking difference between the first four, "bubble" quarters of the average recovery and the following eight quarters. Productivity grows, on average, 2.69 percentage points a year faster than trend in the first four quarters but 0.11 percentage point *slower* than trend over the next eight quarters. This difference is almost

Table 8. The Early-Recovery Productivity Growth Bubble across Seven Cyclical Episodes, 1958–2003

Percent

Episode	Change in productivity deviation from trend	Predicted value	Contribution of EOE effect to predicted value	Predicted value without EOE effect	Residual
First four quarters after real GDP trough					
1958:2–1959:1	3.08	2.43	0.01	2.42	0.65
1961:1–1961:4	3.89	3.90	1.24	2.66	–0.01
1971:1–1971:4	0.67	0.99	2.69	–1.69	–0.32
1975:2–1976:2	2.77	3.57	2.07	1.50	–0.80
1983:1–1983:4	3.22	3.14	2.61	0.53	0.08
1991:2–1992:1	2.43	1.81	1.70	0.11	0.62
2001:4–2002:3[a]	2.79	0.77	1.01	–0.23	2.02
Average	2.69	2.37	1.62	0.76	0.32
Next eight quarters					
1959:2–1961:1	–1.61	–1.17	–0.82	–0.35	–0.44
1962:1–1963:4	0.46	0.01	1.24	–1.23	0.45
1972:1–1973:4	0.72	–0.50	–2.07	1.57	1.21
1976:2–1978:1	–0.40	–0.94	2.07	–3.01	0.55
1984:1–1984:4	0.11	0.98	. . .	. . .	–0.87
1992:2–1994:1	–0.51	0.51	1.70	–1.19	–1.02
2002:4–2003:2[a]	0.46	–0.29	1.01	–1.30	0.75
Average	–0.11	–0.20	0.52	–0.92	0.09

Source: Regression estimates in table 5, column 5-5.

a. Results are based on same regression reestimated with leading values omitted and only the current value of the output deviation included.

entirely captured by our dynamic model of cyclical productivity behavior, with an average residual of only 0.32 percentage point in the first four quarters and an even smaller 0.09 percentage point in the subsequent eight quarters. The predicted change in the first four quarters is 2.37 percentage points compared with –0.20 percentage point in the subsequent eight quarters.

How does the model explain the early-recovery bubble? As shown for the average of the first four quarters, about two-thirds of the predicted 2.37-percentage-point early-recovery productivity growth is attributable to the EOE reversal effect and the remaining one-third to the other variables, almost entirely the role of unusually rapid output growth in the first four quarters in stimulating rapid productivity growth. In the subsequent eight quarters, the EOE reversal effect and the remaining variables work

in opposite directions: the EOE effect continues to boost productivity growth, while slower output growth and the error correction term both hold down productivity growth relative to trend.

This history of cyclical recoveries puts the 2001–03 productivity performance into perspective. Relative to trend, the early take-off of actual productivity growth in the first four quarters after the GDP trough in 2001:3 was not unusual and in fact was roughly equal to the average in the previous six episodes. What was unusual is that the productivity equation could not explain most of this growth, and the residual for the first four quarters was 2.02 percentage points compared with an average of 0.32 percentage point. The estimated EOE effect was smaller in 2001–03, and actual real GDP growth was slower than trend, holding down predicted productivity growth. As measured by the residuals, the initial four quarters (2001:4–2002:3) were more unusual in historical context than the final three quarters. As noted previously, these comments on deviations from trend growth must be evaluated in the context of a robust trend slightly above 3.0 percent a year, which is even faster than the previous peak of trend growth in 1961–62.[34]

Labor Market Counterparts of Rapid Productivity Growth

Economists have marveled at the productivity growth data examined above, but their counterpart through mid-2003, a continual erosion of payroll employment, was a source of despair for many households, not to mention incumbent politicians. A virtue of the output identity used here is that one can use it to trace the unusual strength of productivity growth to unusual weakness in other elements of the identity. Table 9 does this for three subperiods: early 1985 to late 1995, late 1995 to mid-2000, and mid-2000 to mid-2003. All elements of the identity are listed in the same order as in table 3, and the final column sums the right-hand-side variables of equations 2 and 3.

Some of the special aspects of the past three years were already revealed in table 3, which examined the ratio of actual to trend behavior for these same eight variables. Now the statistical evidence presented above can be added in the form of predicted values and residuals, thus

---

34. The trend growth rate of productivity during the four quarters of 1961 was 2.95 percent a year; the fastest rate recorded by my technique was between mid-1948 and mid-1949, when it reached 3.15 percent a year.

Table 9. Annual Growth Rates of Actual and Trend Real GDP and Its Components, Selected Intervals, 1987–2003

Percent

Episode	Real GDP	Output per hour[a]	Hours per employee[b]	Employment rate	LFPR	Working-age population	Mix effect	Payroll-household employment ratio	Sum of last seven columns
1987:1–1995:4									
Actual	2.71	1.45	–0.07	0.13	0.19	1.08	–0.44	0.38	2.71
Trend	2.77	1.59	–0.04	0.04	0.23	1.13	–0.43	0.24	2.74
Ratio of actual to trend	–0.06	–0.14	–0.03	0.09	–0.04	–0.05	–0.01	0.08	–0.03
Predicted ratio	. . .	0.21	–0.04	–0.01	–0.18	. . .	0.06	–0.03	0.12
Residual	. . .	–0.35	–0.07	0.10	0.13	. . .	–0.07	0.11	–0.15
1995:4–2000:2									
Actual	4.20	2.56	–0.04	0.38	0.22	1.29	–0.71	0.49	4.20
Trend	3.28	2.35	–0.17	0.12	0.04	1.24	–0.70	0.37	3.27
Ratio of actual to trend	0.92	0.21	0.13	0.26	0.18	0.05	–0.01	0.21	0.93
Predicted ratio	. . .	0.69	–0.19	0.36	–0.25	. . .	–0.09	0.06	0.54
Residual	. . .	–0.48	0.32	–0.11	0.43	. . .	0.08	0.15	0.39
2000:2–2003:2									
Actual	1.50	3.46	–0.83	–0.78	–0.37	1.21	–0.68	–0.51	1.50
Trend	3.39	3.03	–0.42	0.04	–0.03	1.22	–0.65	0.19	3.38
Ratio of actual to trend	–1.89	0.43	–0.41	–0.82	–0.34	–0.01	–0.03	–0.81	–1.88
Predicted ratio	. . .	–0.75	–0.36	–0.79	–0.08	. . .	0.37	–0.68	–2.19
Residual	. . .	1.18	–0.05	–0.03	–0.26	. . .	0.40	–0.13	0.31

Source: Author's calculations; see appendix A.

a. Predicted values and residuals are from regressions reported in table 5, column 5-5, with four lags on the output deviation variable substituted for four leads. Output per hour data are for the nonfarm private business sector. Predicted values and residuals for other columns (except working-age population) are from the regressions reported in table 4.

b. Output per hour data are for the nonfarm private business sector.

allowing a determination of how much of the observed change in the actual-to-trend ratio for each variable was normal cyclical behavior and how much was surprising in a historical context. To focus first on the bottom panel of table 9, the startling event of the past three years was that real GDP grew 1.89 percentage points a year slower than trend, whereas productivity growth grew 0.43 percentage point a year faster than its own accelerating trend. Balancing this differential was the fact that all the other components of the identity grew slower than trend; in order of size: the employment rate (–0.82 percentage point), the payroll-to-household employment ratio (–0.81 percentage point), hours per employee (–0.41 percentage point), the LFPR (–0.34 percentage point), and the mix effect (–0.03 percentage point).

My econometric analysis balances its large positive residual in explaining productivity growth relative to trend with a large negative residual on the mix effect and smaller negative residuals on the LFPR and the payroll-to-household employment ratio. In trying to understand the recent period, it is helpful to combine the mix and employment ratio effects; as in equation 2, these two effects when multiplied together become the ratio of real GDP per household employee to the ratio of NFPB output to NFPB payroll employees. The sum of the growth rates of these two effects is a change in the actual-to-trend ratio of –0.84 percentage point a year, with a predicted value of –0.31 percentage point and a residual of –0.53 percentage point. An examination of the components of these ratios suggests that the main issue is the behavior of payroll employment rather than the ratio of real GDP to NFPB output.[35]

This section has emphasized the unusual behavior of productivity growth relative to estimated trend productivity growth in the period 2000–03 and its counterpart in slow growth in hours of work. But productivity and hours were not the only unusual aspects of the post-1995 U.S. economy. Notably different from past business cycles was the unusual decline in the rate of capacity utilization from 83.6 percent in 1995 to 82.7 percent in 2000, during a period when the ratio of real GDP to trend

---

35. I have estimated a trend for the ratio Q/Q^B, which is steadily negative at a rate of about –0.4 percent a year over the past decade. This ratio is strongly countercyclical, since NFPB output declines relative to real GDP in recessions and in the early stages of recoveries. The cyclical increase in this ratio in 2001–02 was typical, indicating that the counterpart of unusual productivity behavior was unusual behavior of employment, not of the Q/Q^B ratio.

was increasing from –1.2 to +2.3 percent. Capacity utilization in 2002 was lower than in any year since 1983, despite the much lower ratio of real GDP to trend in the earlier year (–6.2 percent compared with –1.8 percent in 2002). Linking the behavior of capacity utilization with that of productivity is the common theme of rapidly rising capacity to produce ICT investment. I now turn to the substantive puzzles posed by recent productivity behavior, in which ICT investment plays a central role.

Explanations of Rapid Productivity Growth in 2000–03

The second major task of this paper is to examine possible substantive explanations for the post-2000 productivity growth acceleration. The two most compelling hypotheses are, first, that an unusual degree of downward pressure on profits led to unusually aggressive cost cutting by firms, and second, that intangible capital acted as a source of dynamic adjustment in the response of productivity growth to a boom in ICT capital investment. The need to accumulate intangible capital in order to realize the potential of ICT investment can result in a delay of several years in the productivity benefits of that investment. It can also cause mismeasurement of the timing of productivity growth that contributes to explaining some of the post-2000 productivity growth acceleration. After examining these hypotheses, I assess several less important substantive issues, including outsourcing and changes in the structure of labor markets, as well as the hypotheses that payroll employment, hours per employee, or both, have been mismeasured to an unusual extent during the last few years.

Cost Cutting and the Profit Squeeze

Both productivity and profits are leading indicators. Because of lags in hiring and firing, productivity growth leads output, surging when output is growing rapidly and then slowing down as labor input catches up to output in the typical expansion. Late in the expansion, output growth slows, but overly optimistic expectations may lead firms to continue adding to labor input during this late-expansion slow-growth period, thus reducing productivity growth to a smaller positive or even a negative rate. Although some of this typical cyclical dynamic is captured by the adjustment lags in

my econometric equations for productivity and the other components of the output identity, the tendency to overhire in the last stages of the expansion goes beyond the simple dynamics of lagged coefficients on output changes, and this is what is captured by the EOE dummy variables discussed above.

Profits are related to productivity through the income shares of labor and capital. By definition, labor's share is equal to real compensation per hour divided by output per hour. If increases in compensation lag behind productivity in the early phases of a cyclical expansion, labor's share will decline and capital's share will rise, and the rate of return on capital will rise even faster to the extent that the rising utilization of capital causes an increase in capital productivity.[36] The cyclical expansion of the 1990s exhibited behavior typical of corporate profits as measured in the National Income and Product Accounts (NIPA), which registered a near doubling of nominal profits between 1992 and 1997 followed by a decline between mid-1997 and early 2000, a slight further decline into 2001, and then a recovery in 2002 and early 2003 back to the 1997 nominal peak and beyond.[37]

William Nordhaus contrasts the behavior of NIPA profits with that of S&P reported profits, which show a very different timing pattern, growing by 70 percent between early 1998 and early 2000 and then declining by more than half between early 2000 and early 2001.[38] He attributes a substantial role in this "most unusual pattern" to a wide variety of shady accounting tricks to which corporations turned as they desperately attempted to pump up reported profits during 1998–2000 in an environment in which true profits were declining. In Nordhaus's words, these tricks led to the "enrichment of the few and depleted pension plans of the

36. To be explicit, one can write W as the nominal wage rate, A as aggregate hours of work, R as nominal profit per unit of capital, K as the number of units of real capital, P for the GDP deflator, and Q for real GDP. When productivity (Q/A) grows faster than real compensation per hour (W/P), the share of labor (WA/PQ) declines while the share of capital (RK/PQ) rises. If rising capital utilization increases the productivity of capital (Q/K), then by definition the real rate of return (R/P) must rise more rapidly than capital's income share.

37. Nominal corporate profits after tax with inventory valuation and capital consumption adjustments were $309.9 billion in 1992, $596.6 billion in 1997, $528.7 billion in 2000, and $532.3 billion in 2001, recovering to $662.7 billion in 2003:2 (*Economic Report of the President*, February 2003, table B-90, p. 381, column 3, and www.bea.doc.gov/bea/newsrel/gdp203f.xls).

38. Nordhaus (2002a).

many." A further unusual aspect of 2001–02 was the extremely low ratio of S&P reported earnings to S&P operating earnings, primarily due to the one-time charges that firms take to correct for previous business or accounting mistakes. Overall, Nordhaus estimates that reported S&P earnings for 2001 were held down by about 30 percent by a combination of normal cyclical and extraordinary accounting impacts.

The unusual trajectory of S&P reported profits in 1998–2001 placed unusual pressure on corporate managers to cut costs and reduce employment. During the 1990s corporate compensation had shifted to relying substantially on stock options, leading first to the temptation to engage in accounting tricks during 1998–2000 to maintain the momentum of earnings growth, and then sheer desperation to cut costs in response to the post-2000 collapse in reported S&P earnings and in the stock market. The stock market collapse had an independent impact on the pressure for corporate cost cutting, beyond its effect on the stock-option portion of executive compensation, by shifting many corporate-sponsored defined-benefit pension plans from overfunded to underfunded status.

A plausible interpretation of the unusual upsurge of productivity growth in 2002–03, then, is that it was the counterpart of an unusual degree of pressure for corporate cost cutting, in turn caused by the role of accounting scandals and corporate write-offs that led to the unusual trajectory of reported S&P profits relative to NIPA profits. The unusual nature of recent corporate cost cutting has been widely recognized. As the *Wall Street Journal* put it:

> The mildness of the recession masked a ferocious corporate-profits crunch that has many chief executives still slashing jobs and other costs. . . . Many CEOs were so traumatized by last year's profits debacle that they are paring costs rather than planning plant expansions.[39]

This chain of causation from the profits "debacle" to the 2002–03 productivity surge seems plausible as the leading explanation of the unusual productivity behavior documented in previous sections. But it raises a central question: How were corporate managers able to maintain output growth while cutting costs so savagely? Why didn't the massive layoffs cause output to fall, as it would have if productivity growth had stagnated? This

---

39. Jon E. Hilsenrath, "While Economy Lifts, Severe Profit Crunch Haunts Companies; Nervous CEOs Could Slow Recovery by Continuing Layoffs, Plant Closings," *Wall Street Journal,* April 1, 2002, p. A1.

brings us to the central role of ICT investment in the post-1995 productivity growth revival and to the puzzle of explaining how productivity growth surged after 2000 even as ICT investment growth was collapsing along with corporate profits and the stock market.

Why Did Productivity Growth Accelerate after 2000 While ICT Investment Collapsed?

As shown above, productivity growth after mid-2000 was substantially more rapid than during the initial revival period of 1995–2000. Studies of the initial 1995–99 productivity revival period, most notably those of Dale Jorgenson and Kevin Stiroh, and Stephen Oliner and Daniel Sichel,[40] attributed a large fraction of that revival to the *production* and the *use* of ICT equipment and software. In the case of Oliner and Sichel, the analysis included not only ICT capital, including software, but also the semiconductors that powered the hardware. This leads us to the paradox that productivity growth proceeded to a second stage of acceleration during 2000–03 just as the ICT investment boom collapsed.

Data on real investment in computers and other products with rapid relative price changes become increasingly more misleading as time extends past the base year in the NIPA, currently 1996. To avoid potential errors of interpretation, the correct measure of the importance of ICT investment is the *nominal* share of that particular type of investment in *nominal* GDP. For computer hardware, that share averaged 0.96 percent in 1997–2000 but then crashed to 0.71 percent in 2002 and by 2003:2 had recovered only to 0.77 percent. A more comprehensive measure that includes not just computer hardware but also software and "other" (mainly communications) equipment registered an average GDP share of 4.23 percent in 1997–2000, reaching 4.55 percent in 2000, and then fell to 3.83 percent in 2002 before recovering to 3.94 percent in 2003:2.

The Oliner-Sichel Decomposition

The most influential research supporting a large role for ICT investment in the post-1995 productivity growth revival appears in a series of papers by Oliner and Sichel (O-S).[41] Their approach, presented in

---

40. Jorgenson and Stiroh (2000); Oliner and Sichel (2000).
41. Oliner and Sichel (2000, 2002).

table 10, has attracted wide attention because of the clarity with which they distinguish the role of deepening of ICT capital (third line of table 10), that is, the benefits of rapid ICT investment to the *users* of ICT capital, from that of the *production* of ICT capital in raising the growth rate of total factor productivity (TFP) for the economy as a whole (tenth line). It is easy to follow the changes in the O-S results as new data emerge, because they are always presented in the same format, and because the same initial time period (growth rates from 1973 to 1995) is compared with the revival period (1995 to the latest year for which data are available).

Table 10 compares the initial (2000) O-S decomposition with their latest unpublished results, which extend the findings to 2002.[42] The table first reports the growth rate of labor productivity and then subtracts the contributions of capital deepening and improvements in labor quality (education) to arrive at the growth rate of TFP. The capital deepening component is further subdivided into ICT and other capital and into three types of ICT capital. The resulting TFP growth rate is then decomposed into the role of the *production* of ICT capital and that of all other contributions to TFP growth.

Finally, table 10 adds two lines to the standard O-S decomposition. The two types of ICT capital contribution, ICT capital deepening and the TFP effect, can be added together, as is done in the next to last line of the table. Then the total ICT contribution can be divided by the growth rate of labor productivity, from the first line of the table, to yield the total contribution of ICT capital to productivity growth and to the productivity revival, as shown in the last line of the table. Table 10 shows three different decompositions of the post-1995 productivity growth revival. The first, using data through 1999 (first three columns), is taken from the initial O-S paper, and the second uses the latest (revised) data for the same period (second three columns). Data revisions reduce the overall productivity revival while leaving the ICT contribution intact, and this boosts the contribution of ICT capital to the 1995–99 revival from 81 percent to 98 percent.

The third decomposition extends the endpoint of the data from 1999 to 2002 (last two columns of table 10). The revival in the growth rate of labor productivity increases from 0.96 to 1.20 percentage points, while

---

42. There have been intervening analyses of data ending in 2000 and 2001 (see Oliner and Sichel, 2002); these are omitted to simplify the table and because the most recent results are of greatest interest.

Table 10. Decomposition by Oliner and Sichel of Growth in Labor Productivity, 1973–2002

Percent

Component	Original O-S estimates (2000)			Latest O-S estimates (2003)				
	1973–95	1995–99	Change, 1973–95 to 1995–99	1973–95	1995–99	Change, 1973–95 to 1995–99	1995–2002	Change, 1973–95 to 1995–2002
Labor productivity growth	1.41	2.57	1.16	1.41	2.36	0.96	2.61	1.20
Contributions from:								
Capital deepening	0.77	1.10	0.33	0.72	0.98	0.26	1.20	0.49
ICT capital	0.46	0.96	0.50	0.42	0.95	0.53	0.93	0.51
Computer hardware	0.25	0.59	0.34	0.22	0.52	0.31	0.47	0.25
Software	0.12	0.27	0.15	0.12	0.33	0.20	0.33	0.21
Communications equipment	0.08	0.10	0.02	0.08	0.09	0.01	0.13	0.05
Other capital	0.31	0.14	–0.17	0.30	0.03	–0.26	0.27	–0.02
Labor quality	0.27	0.31	0.04	0.27	0.30	0.03	0.25	–0.02
Total factor productivity	0.36	1.16	0.80	0.42	0.98	0.56	1.15	0.74
Of which contributed by:								
ICT and semiconductors	0.22	0.66	0.44	0.30	0.72	0.41	0.70	0.40
Other nonfarm business	0.15	0.50	0.35	0.11	0.26	0.15	0.45	0.34
Memorandum: Total contribution of ICT and semiconductors[a]	0.67	1.62	0.95	0.72	1.66	0.94	1.63	0.91
As share of total change in labor productivity	47.93	63.04	81.29	51.55	70.37	98.01	62.62	75.53

Sources: Oliner and Sichel (2000, tables 2 and 4); data provided by Daniel Sichel, Federal Reserve Board, to the author.
a. Sum of contributions of ICT capital deepening and ICT and semiconductor contributions to TFP.

the total contribution of ICT capital shows surprisingly little response to the decline in ICT investment discussed above. As a result, the contribution of ICT capital declines from 98 percent in the period ending in 1999 to 76 percent in the period ending in 2002.[43] The spurt in productivity growth from 1999 to 2002 is more than explained by capital deepening in "other capital" and more rapid TFP growth contributed by sectors of the economy other than ICT and semiconductor capital. The puzzling absence of a decline in the ICT contribution as well as the upsurge in "other" capital deepening can both be traced to the same cause: the rapid decline in hours of labor input in 2001–02. Since all the capital deepening terms in table 10 represent the change in a capital-labor ratio times an income share of that type of capital, the apparent resilience of the ICT role disguises the fact that the ICT contribution by itself fell by half between 2000 and 2002, but this is dampened by the rapid decline in labor input.

Delay and Intangible ("Hidden") Capital

Drawing back from the details of table 10, one can take a broader view of the claim that, at least through 1999, virtually all of the productivity growth revival can be attributed to the production and use of ICT capital. This finding seems compatible with numerous studies, especially those of Jack Triplett and Barry Bosworth and of Nordhaus,[44] which pinpoint wholesale and retail trade and securities trading as the industries outside of ICT manufacturing where the productivity growth revival is most evident. In these analyses the largest single contribution to the revival is deepening in ICT capital, and this is precisely the capital that has been used so effectively in trade and securities trading.

Nevertheless, several questions may be raised about the implication of table 10 that the post-1995 productivity revival, at least through 1999, resulted entirely from increased production and use of ICT equipment.[45] The O-S technique requires that the full productivity payoff from the use

43. Data revisions released on August 7, 2003, would further reduce the ICT share of the 1995–2002 revival from 76 percent to 67 percent, allowing only for the revisions in labor productivity and assuming no revisions for any other figure in the final column of table 9.

44. Triplett and Bosworth (2002); Nordhaus (2002b).

45. One of the interim versions of the O-S decomposition went through an endpoint of 2001 and concluded that ICT capital *overexplained* the revival when dated as 1995–2001 (see Sichel's discussion of Nordhaus, 2002b).

of computers occur at the exact moment that the computer is produced.[46] Assuming no delay between production and installation, the computer produces its ultimate productivity benefit on the first day of use. Numerous observers, led by Paul David,[47] argue instead that there is in fact a substantial delay in reorganizing business practices to take advantage of new hardware and software. If the O-S method fails to take such delay into account, it will exaggerate the contribution of ICT capital deepening to the post-1995 revival during the years of peak ICT investment. Then, in the period 2001–03 when ICT investment has declined, it will understate the contribution of ICT capital by omitting the delayed benefits from previous ICT investment. This would lead the method to understate the acceleration of TFP growth in the years of peak ICT investment and to overstate it during the period of declining ICT investment.

David's delay hypothesis was based on a very general analogy between the invention of electricity and computers. Since his contribution was written, electricity, the internal combustion engine, and computers—the three major inventions of the last 125 years that have wide and varied uses across industries—have been christened "general purpose technologies."[48] They have in common that they require complementary inventions and business reorganization to become fully effective. It took forty years for electricity to make a major impact on manufacturing productivity, because factories had to be reorganized and machines reduced in size to bring power to the individual workstation. The internal combustion engine required massive investment in highways and airports to reach its full productivity-raising potential. The role of business reorganization and process improvement in the form of "intangible capital" that is complementary to ICT investment has been the focus of recent interpretations of the post-1995 productivity growth revival by Susanto Basu and coauthors (Basu, Fernald, Oulton, and Srinivasan, B-F-O-S) and by Shinkyu Yang and Erik Brynjolfsson (Y-B).[49] Both groups of authors argue that measured investments in computer hardware and software

46. Recall that the GDP statistics on which they rely measure output by production and treat any unsold goods as inventory accumulation, a part of GDP.
47. David (1990).
48. Bresnahan and Trajtenberg (1995).
49. Basu and others (forthcoming); Yang and Brynjolfsson (2001). A precursor to the B-F-O-S paper is Basu, Fernald, and Shapiro (2001). A follow-up to the Y-B paper is Brynjolfsson, Hitt, and Yang (2002). An early version of the Y-B hypothesis is presented in Brynjolfsson and Yang (1997).

require complementary, unmeasured investments in intangible capital, including business reorganization, new business processes, retraining, and general acquisition of human capital. The direction of mismeasurement is the same as implied by the delay hypothesis for the O-S estimates: it overstates the extent of the productivity growth revival between 1995 and 2000 but understates it afterward.

The intangible capital hypothesis is intriguing, because it offers a possible explanation of the puzzling second acceleration of productivity growth after 2000, in the wake of the collapse in the ICT investment boom. Intangible capital is complementary to ICT and contributes to measured output. However, net intangible investment, which builds intangible capital, leads to an understatement of current-period output, because it is not counted in measured output even though the factor inputs required to produce it are counted. When the share of intangible investment is constant, this means that the amount of output credited to other, tangible factors and to total factor productivity is overstated by the percentage of output devoted to intangible investment, but there is little effect on growth rates or on the allocation of growth between factor inputs and TFP. However, when the share of output devoted to intangible investment is changing, the effect on measured productivity and its growth rate can be dramatic. A 1-percentage-point increase in the share of output devoted to intangible investment, say, from 3 to 4 percent, reduces measured output relative to total output almost point for point. If this increase takes one year, measured output growth is biased downward by 1 percentage point for one year; if it takes five years, the growth rate bias is one-fifth as large, or 0.2 percent a year.

How large might these hypothesized measurement effects be? As a rough approximation, assume that intangible ICT investment is of about the same magnitude as measured ICT investment. The ratio of total measured ICT investment to nominal GDP rose from 3.54 percent in 1995:3 to 4.59 percent in 2000:3 before falling back by about half its previous increase, to 4.01 percent in 2003:3. For a narrower concept of ICT investment, one that includes only computer hardware and software and excludes communications equipment, photocopiers, and other office equipment, the investment shares in these same periods were 2.03, 2.81, and 2.57 percent. (It is interesting that this burst of growth in the ICT share in the late 1990s was not unprecedented: the share of total ICT investment increased from 2.01 percent of nominal GDP to 2.98 percent

between 1978:1 and 1983:4.) To calculate the effects on the measurement of output and productivity, however, we need the growth in real rather than nominal investment. Over 1995–2000, tangible real ICT investment grew at a 17.7 percent average annual rate, or 13.7 percentage points faster than real GDP. With a 4 percent average share of ICT in GDP during the period, this implies that the (unmeasured) share of real intangible investment was increasing by 0.5 to 0.6 percentage point a year, and hence annual growth of GDP and productivity were understated by the same amount. Using the narrower definition of ICT investment, with its annual growth rate of 22.5 percent and average nominal GDP share of 2.4 percent, gives a similar estimate. During 2000–03 this understatement of output and productivity would have been reversed, although only in part, since an understatement would be expected for the entire 1995–2003 period when the ICT share rose. Thus the hypothesized effects, even if realistic, could account for only a part of the surprising recent strength in productivity. It is worth noting that these effects are not evidence for or against various hypotheses about delays in realizing the benefits of ICT investment, which simply assert that the contribution to output is not fully realized at the time the measured investment is made.

The Y-B hypothesis has implications for earlier years as well. The share of spending on computer hardware was growing rapidly, particularly during 1972–87, the interval that led Robert Solow to utter his famous quip, which later became known as the Solow paradox, that "we see the computer age everywhere except in the productivity statistics." Going beyond the specific restrictions imposed by particular models, the role of delayed benefits from the rapid growth in ICT investment in the late 1990s seems incontrovertible. Jeffrey R. Immelt, chief executive officer of General Electric, refers to the delayed benefits of ICT spending by saying, "It takes one, two, three years to get down the learning curve and figure out new ways to use it." Cisco CEO John Chambers estimates the learning curve at more like five to seven years.[50]

50. Quotes from Peter Coy with Timothy J. Mullaney, "Still Getting Stronger," *Business Week,* September 15, 2003, pp. 32–35. Brynjolfsson and Hitt (forthcoming) find evidence of a delay in computer effects on firm-level productivity, and Bresnahan, Brynjolfsson, and Hitt (2002) find direct evidence of large, time-consuming organizational investments that complement computer investments on production functions.

Other Substantive Explanations

In this section I consider briefly two other suggestions regarding structural changes over the past few years that could have explained some of the upsurge in productivity growth. These include outsourcing and several changes in labor markets that may increase flexibility and reduce the barriers to adjusting the size of the work force in response to a profit squeeze.

OUTSOURCING. The United States has been outsourcing jobs on a large scale for the past two decades in the form of its ongoing trade deficit. The claim that outsourcing is something new, and related to the post-2000 productivity growth upsurge, refers to white-collar jobs that American firms are shifting to call centers and offshore programming specialty establishments in countries like India. Consider a firm initially producing 1,000 units of output with 100 workers, for a productivity level of 10 units per worker. Assume that total labor cost is 800 and profits are 200. Now imagine that 10 jobs are shifted to India and the new Indian workers are paid one-quarter of the American wage, or 20 units per 10 jobs instead of the previous 80 units. The firm must pay 20 for the imported services, reducing output from 1,000 to 980 and raising productivity from the initial 10 to 980 divided by 90, or 10.9. Productivity has increased by 9 percent, and profits have increased from 200 to 280, or by 40 percent.

It is unlikely that the entire benefit of the cost savings through outsourcing will be retained by the firm. Competitive pressure will reduce prices, and properly computed deflators will maintain the original level of real output minus the spending on imported services. Thus, in principle, outsourcing does increase U.S. productivity, but the question is by how much. Anecdotal evidence suggests that white-collar jobs outsourced over the past three years number in the hundreds of thousands rather than millions, suggesting a productivity impact of a few basis points rather than a few percentage points.[51] Clearly, however, outsourcing of call center and programming jobs is a by-product of ICT innovation in the late

51. Supporting the view that the number of outsourced jobs is in the hundreds of thousands rather than millions, Goldman Sachs estimates that the number of services sector jobs outsourced by U.S. companies to their foreign affiliates is 200,000 (Michael Schroeder, "India Aims to Calm U.S. Outsourcing Fears," *Wall Street Journal,* November 13, 2003, p. A4). This does not include jobs outsourced to foreign-owned contractors, which might double that number.

1990s, particularly innovation that drastically reduced the price of always-on broadband communications.

LABOR MARKET FLEXIBILITY. That there is a set of institutions that make American labor markets more flexible than those in some other developed countries, particularly in Western Europe, is likewise nothing new. The contrast between labor market institutions in the United States and in Europe has been on the front burner of comparative macroeconomics since the mid-1980s, when economists on both sides of the Atlantic started trying to explain why European unemployment had increased from below the American level in the 1960s to above it after 1985.[52] The question is whether the past several years have seen a significant further movement in the direction of flexibility in the U.S. labor market that could have facilitated corporate cost cutting and the achievement of faster productivity growth.

A review of several labor market indicators casts doubt on the flexibility hypothesis, simply because these indicators exhibit little change in recent years. The share of the labor force working part-time has fluctuated in a narrow range around 16 percent since 1984. The Bureau of Labor Statistics (BLS) series on the number of self-employed was roughly 9.4 million in mid-2003, compared with 9.3 million in early 2000. The erosion of union membership has remained on a steady downward glide path since 1984 after a sharp ratchet downward in 1980–83. The increased role of temporary help agencies and other types of contingent worker contracts was already evident to labor economists in 1999, before the recent productivity growth take-off.[53]

Mismeasurement of Hours per Employee and Payroll Employment

Moving beyond substantive hypotheses, at least two possible dimensions of mismeasurement are independent of the intangible capital mismeasurement hypothesis discussed above. These seek to explain some fraction of the post-2000 upsurge in productivity growth by arguing, for different reasons, that hours per employee and payroll employment are both understated, implying that aggregate hours of work are understated and productivity is correspondingly overstated. Two of the counterparts

---

52. See, for instance, Lawrence and Schultze (1987).
53. See Katz and Krueger (1999, especially pp. 45–60).

to the rapid productivity growth in 2000–03 documented in table 9 are the unusually rapid decline in hours per employee, occurring at an annual rate of –1.16 percent, and the decline in the ratio of payroll to household employment, at an annual rate of –0.51 percent.

HOURS PER EMPLOYEE. Among others, Stephen Roach of Morgan Stanley has long argued that hours per employee are understated because modern technology, especially cellular telephones, laptop computers, and at-home desktop computers, allows work to be performed outside of the formal workday, and employers are under no restraint in expecting their employees to put in extra hours. So what else is new? Many professionals, including novelists and even economics professors, have been writing late into the night since the invention of the quill and candle, much less the typewriter or the personal computer. Granted, employer expectations of extra hours of work now extend to a broad professional class of lawyers, consultants, investment bankers, and others who formerly had to do most or all of their work within the confines of the office during the business day.

Yet one can be skeptical for several reasons. First, even before the invention of the personal computer, work could be taken home in the form of paper: lawyers read and revised briefs, doctors read medical journals, and even the reading of business magazines counts as work for consultants and investment bankers. Second, there was no sudden discontinuity in 2000 or 2001 in hours worked at home or out of the office; if hours of work are undercounted, the source of the undercounting has grown gradually in importance over the past two decades. Third, some unknown fraction of work outside the office is offset by consumption inside the office. The mix of work at home and consumption on the job has become so entangled that it seems hopeless to think of any feasible measurement project that could capture the extent of mismeasurement in all its nuances. As a recent *Wall Street Journal* article explains,

> Between meetings and memo writing, employees and their bosses log onto the Internet to pay bills, manage their 401(k) accounts, shop for birthday gifts and plan vacations. While attending their children's school events, they use their BlackBerrys to answer colleagues' questions, proof letters, and make appointments with clients.[54]

54. Carol Hymovitz, "More Managers Allow Workers to Multitask As Job and Home Blur," *Wall Street Journal,* October 28, 2003, p. B1.

PAYROLL EMPLOYMENT. An unusual feature of the 2001–03 period was a marked decline in employment as measured by the payroll survey relative to employment as measured by the household survey (referred to above as the payroll-to-household employment ratio, or E^p/E). Recently John Kitchen has developed the intriguing argument that payroll employment has been significantly understated during the period since its last benchmark revision in March 2002.[55] Kitchen argues that, early in a recovery, payroll employment typically shows greater weakness than household employment, but that it is typically revised upward when such discrepancies arise. To demonstrate this, Kitchen compares "current vintage" payroll employment data—for example, 1983 data as recorded in 2003—with contemporary data available in 1983 and 1984 government publications. To demonstrate that revisions in payroll data are influenced by contemporary household data, he runs a regression over three early-recovery periods (1975–76, 1983–84, and 1991–92) in which the current vintage payroll data are explained by contemporary payroll data and contemporary household data. The fact that the household data enter with a significant positive coefficient implies that the payroll data are revised upward by a larger amount, the greater the positive discrepancy between household and contemporary payroll data.

Applying his estimated regression equation to the current period, Kitchen estimates that, in July 2003, payroll employment was understated by 570,000, or about 0.44 percent. Translating this directly into an overstatement in the productivity data would reduce output per hour in the NFPB sector in 2003:2 by 0.44 percent and would reduce its annual growth rate between 2000:2 and 2003:2 from 3.46 to 3.31 percent, a relatively small downward revision.

However, Kitchen's plausible argument has been contradicted by the BLS's announcement, in early October 2003, that payroll jobs for March 2003 were being revised *downward* by 145,000, or about 0.1 percent of the total, not upward as Kitchen had predicted. That small downward revision would apply through 2003:1, and levels thereafter would presumably be affected as well. The revised productivity estimates that incorporate the payroll revision will not be released until February 2004. Kitchen still expects, based on the precedent of 1992–93, that the revision to payroll employment to be announced in October 2004 for the period

55. Kitchen (2003).

beginning in 2003:1 will be a substantial upward revision; the 1992 benchmark revision was a small negative, but the 1993 benchmark was a significant upward revision that resulted in higher growth in payroll jobs for the period from March 1992 to March 1993. He feels that "the jury is still out" about payroll jobs growth after March 2003.[56]

This completes my treatment of the second main goal of the paper, which was to examine substantive and measurement hypotheses that might explain the upsurge of productivity growth in 2000–03—both the more rapid estimated trend growth rate and the positive deviations from trend that are outside the range of historical experience. One can draw from this review of explanations several alternative implications for the future, ranging from the optimistic assessment that the recent acceleration is permanent, to the opposite view that it is inherently temporary. An optimist would point to the estimated increase in trend productivity growth to a rate above 3 percent, faster even than in the peak years of the early 1960s, and to a continuing wave of innovation in creating this acceleration in trend growth. An intermediate view would emphasize the role of intangible investment in spreading out and delaying the productivity benefits of the investment boom of the late 1990s. Although the intangible mismeasurement hypothesis implies that productivity growth may have been overstated during the period of slower ICT investment growth during 2000–03, nothing in this view rules out a new wave of ICT investment and another period like the late 1990s, when productivity growth was understated. The most pessimistic approach emphasizes the unprecedented timing pattern of S&P reported profits, which exhibited an artificial boom in 1998–2000 followed by an utter collapse that let loose a wave of savage cost cutting that continues to this day. Now that costs have been cut, the resulting productivity gains are unprecedented, and profits are resurging to above their 1997 peak. The pressure for cost cutting and hence that unique prop to productivity growth must inevitably fade, but at some unknowable rate.

Connecting the Past to the Future

The third goal of this paper is to forecast growth in potential real GDP over the next two decades through 2023. This requires going beyond a

56. This paragraph is based on a personal communication from John Kitchen.

discussion of trend productivity growth performance to examine the determinants of future growth in the working-age population, in hours per employee, and in the mix and payroll-to-household employment ratio terms that are needed to translate productivity growth in the NFPB sector into an equivalent growth rate for real GDP in the economy as a whole. Although, in principle, future economic growth also depends on changes in the employment rate and the LFPR, likely changes in these ratios on average over the next two decades are too small to make any appreciable difference to potential real GDP growth.[57]

Future Productivity Growth

This paper's extensive review of the U.S. productivity data poses a dilemma in thinking about the future. What growth rate of productivity should be projected over the next twenty years? The estimated productivity growth trend reaches 3.05 percent a year in 2003:2, which is substantially more rapid than the consensus in recent research that the trend has been running at roughly 2.25 percent. In contrast, if one were to base a forecast for the next two decades on the two decades ending in 2003:2, the forecast value would be the actual average over 1983–2003 of 2.03 percent a year. To return to the question posed at the outset, how much of the past is relevant in making future projections?

The estimated trend is, of course, merely a reflection of the past actual outcome and is capable of moving quite rapidly when confronted with a sharp difference between the actual outcome and the current trend. For instance, my own estimate of the productivity growth trend (bottom panel of figure 1) has increased from 1.64 percent in 1994:1 to 2.76 percent in 2000:1, a span of merely six years. The same could happen in reverse if at some point in the future there were a period of years with slow productivity growth.

Several factors lead me to conclude that productivity growth over the next twenty years will be faster than the 2.03 percent annual average registered between 1983 and 2003, not to mention the 1.42 percent average between 1972 and 1995. Productivity growth was held down in the 1970s and early 1980s by the dislocations caused by oil shocks, high inflation

---

57. If the employment rate in 2023 is 94.5 percent rather than 95.0 percent, for example, this implies a reduction in the growth rate of potential real GDP of a mere 0.03 percentage point a year.

and unemployment, and high and volatile interest rates. Numerous econo-
mists have identified a shift toward greater macroeconomic stability in the
mid-1980s. Another reason for optimism is that the 1990s brought a new
and fruitful interaction among innovations in computer hardware, soft-
ware, and communications equipment that appears to have allowed a
greater degree of business reorganization and a greater opportunity for
intangible capital investment than in the 1970s and 1980s.

Working in the opposite direction are several reasons to believe that
productivity growth over the next two decades will be slower than the
mid-2003 estimated trend of 3.05 percent a year. Most important is the
role of the 2000–01 profit squeeze in motivating an unprecedented wave
of cost cuts and reductions in labor input. Some of the productivity boom
has been the counterpart of this cost slashing, and productivity growth is
likely to decline in response to the rapid growth in profits that is already
evident. Second, the intangible capital hypothesis argues that productivity
growth was overstated in the period of slow ICT investment growth in
2000–03, just as it was understated before 2000. Third, many of the most
fruitful applications of the marriage of computer hardware, software, and
communications technology have already occurred. The size of human
fingers and the ability of human eyes to absorb information from tiny
screens set limits to miniaturization. It seems quite likely that diminishing
returns will set in, at some point over the next two decades, in the fruitful
application of the innovation wave of the late 1990s.

Fourth, the recent work of Jorgenson, Mun Ho, and Stiroh emphasizes
that educational attainment in the United States has reached a plateau,
implying that the contribution of changes in labor quality to TFP growth
will taper down gradually to zero over the next two decades from its aver-
age over 1973–2003.[58] Specifically, they estimate the rate of increase of
labor quality to have been 0.38 percent a year in 1995–2001, and they
project that it will decline to a much slower 0.16 percent a year over
2001–11 and fall to a negligible 0.02 percent in 2011–21. When multi-
plied by a labor income share of, say, 0.65, this reduction in labor quality
growth would result in a reduction in trend productivity growth of
0.19 percentage point a year when 2001–21 is compared with
1995–2001. This mechanical conversion of educational attainment into a
measure of labor quality is controversial, and one might argue that the

---

58. Jorgenson, Ho, and Stiroh (2002).

new generation emerging from secondary school with expertise in standard computer software far greater than their parents' might offset most or even all of the plateauing in educational attainment.

Weighing these factors together, I view future productivity growth over the next two decades as more likely to fall in the range of 2.2 to 2.8 percent a year than the 2003 estimated trend of 3.05 percent or the actual 2000–03 growth rate of 3.46 percent. In the concluding section of this paper, I take the mean of 2.2 and 2.8, that is, 2.5 percent a year, as my central projection of future productivity growth in the NFPB sector.

Population Growth

When the trustees of the Social Security Administration forecast potential GDP growth of only 1.8 percent in the period 2040–80, the most important reason for their surprisingly low growth rate is their projection of population growth at only 0.2 percent a year. For comparison, table 1 above shows that growth in the working-age population in the United States over the period 1987–2003 was 1.2 percent a year. Population projections require the determination of three essential numbers, which one may call the "demography triplets": the fertility rate, the mortality rate, and the number of net immigrants. This section discusses each briefly and is based in part on the 2003 report of Social Security's Technical Panel on Assumptions and Methods (TPAM), which is convened every four years to make recommendations on demographic and economic assumptions to the Social Security Advisory Board. All data in this section on population come from the 2003 Social Security Trustees' Report and the 2003 TPAM report.[59]

The starting point for demographic projections is the fact that the replacement rate for fertility, that is, the total fertility rate (TFR) that will maintain a constant population, is 2.1 births per female of childbearing age. In recent years the United States has been running a TFR of 2.05, very close to the replacement rate, in stark contrast to most other rich countries in the world. TFRs in Italy, Spain, Japan, and several other countries have dropped to 1.3, while France and Germany have TFRs of 1.5 to 1.6. Will American exceptionalism in this regard continue, or will the American fertility rate trend down to converge with the rates now prevalent in the European Union and Japan?

59. Board of Trustees (2003); Technical Panel on Assumptions and Methods (2003).

The consensus among demographers seems to favor a continuation of the U.S. TFR in the range of 1.9 to 2.0, slightly less than the recent rate of 2.05 but substantially higher than in Europe. This projection should be viewed in the context of a wide range in the TFR over the postwar period, from a low of 1.74 in 1976 to a high of 3.68 in 1957, the height of the baby boom. When considering how much of the past is relevant for the future, demographers with near unanimity reject any relevance of the baby-boom period of 1947–64 and instead focus on reasons why the U.S. TFR is not likely to decline toward the values observed in Western Europe or Japan.

The first argument in favor of continuity is that the composition of the U.S. population favors relatively high fertility, because of the growing share in the population of Hispanics, who had a TFR in 2001 of 2.75. Even among non-Hispanic white women, however, the TFR during the decade 1991–2002 fell within a range from 1.77 to 1.87, considerably higher than in most other high-income countries.[60] The second argument is more qualitative: women in the United States are better able to combine work and childbearing. The American labor market allows more flexible work hours and makes it easier to leave and then reenter the labor force. The relevance of labor market institutions is validated by a positive correlation between measures of female labor force activity and the fertility rate observed across highly developed nations.[61]

More generally, the higher TFR in the United States results from a set of differences in culture and institutions, including not just the flexibility of labor markets but also the greater independence of young people at an earlier age. Edmund Phelps has provided an economist's description of this difference in youth cultures.[62] American teenagers gain early contact with money and incentives. They start while in high school, working at McDonalds and Burger King. Unlike European college students, who enjoy blanket tuition and stipend grants, most American college students must work for part of their college expenses, borrowing the rest and thereby undertaking an obligation to repay the loans by taking jobs soon after college. Data indicate a much higher incidence of long-term unemployment among European youth than in the United States, and lower

---

60. For basic data on the TFR by demographic group, see Hamilton, Sutton, and Ventura (2003).
61. Brewster and Rindfuss (2000).
62. Phelps (2003).

labor force participation rates. This has an implication for demographic outcomes, because young people in countries like Italy are much less likely to be financially independent at age twenty-five or thirty and more likely to be living with their parents—an atmosphere quite uncongenial to marriage, much less childbearing. These transatlantic differences in labor markets and youth culture may help to explain differences in fertility rates and support the apparent consensus among demographers that American exceptionalism is likely to continue.

Long-term population growth projections also require estimates of mortality rates, which have declined continuously over the last two centuries, but at a highly variable pace. The central question about which opinions vary is the future rate of decline of the age- and sex-adjusted death rate for persons aged 65 and over. That rate declined at an annual rate of 0.78 percent between 1900 and 2000 and 0.91 percent during 1950–2000, but at much slower rates of 0.55 percent between 1979 and 2000 and 0.27 percent between 1996 and 2000. The assumption of the Social Security trustees is an annual rate of decline of 0.71 percent a year between 2002 and 2027. The rate of decline was much slower after 1982 than in the unusual period 1968–82, a fact that some experts attribute to the initial influence of Medicare in bringing down death rates, as free basic medical care was brought to the elderly population for the first time.[63] The 2003 TPAM report recommends projecting a rate of decline over the next seventy-five years equal to that of the last fifty years, that is, 0.84 percent a year (this differs from the historical 0.91 percent rate of decline cited above for 1950–2000 because of adjustments for changes in the relative weight of subgroups within the sixty-five-and-over age group). Their most compelling argument is ironic, in that the TPAM report chose to ignore foreign experience in projecting the fertility rate but finds it relevant in projecting the mortality rate. The report notes that the slow post-1980 rate of decline of mortality in the United States is unusual, and that the more rapid rates of decline in other high-income countries provide a better guide for the future.

The third of the demography triplets is the assumption about future immigration. This is by far the most controversial issue in assessing the official Social Security assumptions about future population growth. The

---

63. Data and this speculation about 1968–82 were provided by Steve Goss, chief actuary of the Social Security Administration.
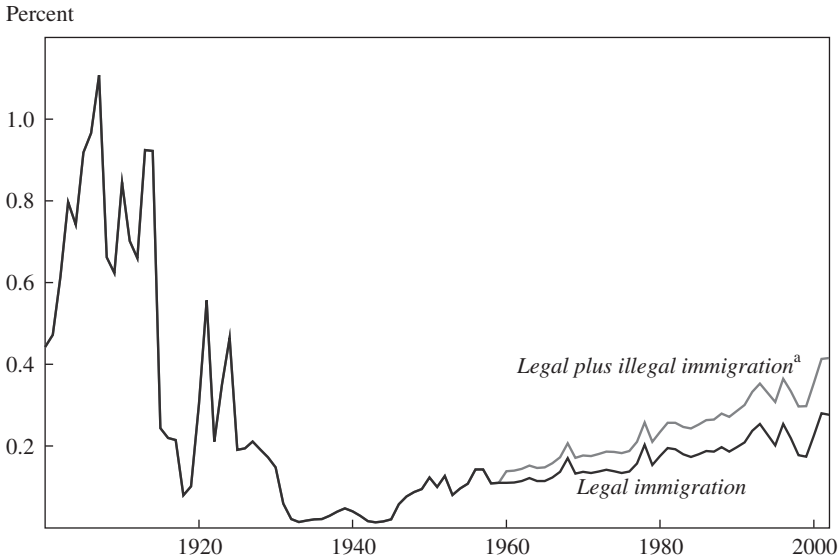
historical context of immigration rates is provided in figure 5, which shows, for the full period 1900–2002, net legal immigration and, after 1960, an estimate of illegal immigration as percentages of the U.S. population.[64] After fluctuating between 0.6 and 1.0 percent of the population most of the time between 1900 and 1913, immigration fell off to a bare trickle between 1933 and 1947, followed by a gradual but steady recovery to 0.41 percent in 2002.

Somewhat incredibly given this historical perspective, which shows a growth rate of total immigration since 1970 of 3.75 percent a year, the Social Security projections assume not just zero growth in the future, but an absolute decline from 1.4 million total immigrants in 2002 to a steady state after 2020 of 900,000 new immigrants a year. This can be translated into an assumption that total net immigration will decline from 0.41 percent of the population in 2002 to 0.21 percent in 2080.

The future behavior of immigration is truly the wild card in projecting future growth in population and hence in potential output. A continuation of the past 3.75 percent annual growth rate of immigration for the next seventy-five years would imply unprecedented ratios of foreign-born to the total U.S. population by the end of that period. But to project negative growth implies a discontinuity that is without any basis. Labor economists would presumably model the immigration process as the intersection of supply and demand, both of which would be likely to increase in the future as the populations of the immigration-supplying countries (Mexico, Central America, China, India, Pakistan, and others) continue to grow, while a growing population and rising income per capita within the United States increase the demand for both skilled and unskilled immigrant labor.

Although many different assumptions could be justified, a conservative basis for estimating future population growth is to forecast that the annual growth rate of immigration will gradually slow from the post-1970 rate of 3.75 percent a year to a rate after two decades (that is, after 2023) that is consistent with a constant ratio of immigration to the population. A simu-

64. Data on population and total legal immigration, excluding immigrants converted from illegal to legal status under the 1986 Immigration Control and Reform Act, were provided by John Wilmoth of the Demography Department at the University of California, Berkeley. Total legal immigration was multiplied by 0.75 to allow for emigration, the same assumption made by the Board of Trustees (2003, p. 81, note 3) at Social Security. Illegal immigration in 2002 was taken to be 400,000, from Board of Trustees (2003, table V.A.1, p. 80) and was assumed to rise linearly from an initial value of 50,000 in 1960.

Figure 5. Net Annual Immigration as a Share of U.S. Population, 1900–2002

Percent

Legal plus illegal immigration[a]

Legal immigration

Sources: U.S. Bureau of the Census and author's calculations.
a. Estimates for illegal immigration are not available before 1960.

lation run on the Social Security Advisory Board's population model indicates that a 1 percent long-run growth rate of both immigration and the population would imply that the ratio of immigration to the population eventually settles down at about 0.5 percent, above the ratio reached in 2001–02 but substantially below the historical peaks reached in 1900–13.

The 2003 TPAM report chose a more conservative approach, assuming that the annual number of net immigrants should increase at half the rate of population growth until the ratio of net immigration to the population levels off at 0.32 percent, the historical average from 1821 to 2002. I find this reliance on 180 years of history to be incongruous with every other demographic and economic assumption in the report and see no reason why the annual growth rate of immigration should suddenly drop today from its recent average of 3.75 percent to a mere 0.6 percent. The tapered decline in the growth rate assumed here makes only a small difference to projected population growth over the next twenty years; it makes a much bigger difference over the following fifty-five years.

Overall, my projection of population growth for 2003–23 is 0.98 percent a year, somewhat higher than the projection for the same period in the 2003 TPAM report. The difference is due entirely to the more gradual tapering off of the growth rate of immigration.[65]

Hours per Employee

The last crucial number in estimating future potential GDP growth is hours per employee. Discussion of hours per employee can be confusing, because there are separate measures of hours per employee in the NFPB sector and of hours per household employee in the total economy. It is this latter concept that the 2003 TPAM report projected to grow at a rate of zero over the next seventy-five years. But the growth rate of NFPB hours per employee has a strong negative trend. How can we reconcile these apparently conflicting measures?

The usual concept of hours per employee, discussed previously, is calculated as aggregate hours in the NFPB sector divided by payroll employment in that sector. As shown in table 11, the actual 1987–2001 growth rate of NFPB hours per employee was –0.15 percent a year, representing the difference between growth in NFPB aggregate hours of 1.46 percent a year and growth in NFPB employment of 1.61 percent a year. Between 2001 and 2003 the decline in hours per employee proceeded at a faster rate of decline of –0.45 percent a year, representing a sharp downturn in the estimated trend with little change in the ratio of actual to trend (see table 1).

Turning to the aggregate economy in the first column of table 11, I assume that hours per payroll employee fell at the same modest rate during 1987–2001 as in the NFPB sector. Because payroll employment grew 0.41 percentage point a year faster than household employment during 1987–2001, hours per household employee *rose* at an annual rate of 0.26 percent. This surprising fact represents the offsetting impact of two long-term trends. More household members today hold multiple jobs,

65. The difference of 0.11 percentage point was calculated in a simulation of the Social Security population model intended to predict the effects of a wide range of alternative immigration assumptions, including the assumption recommended here and the slower-growth assumption endorsed by the 2003 TPAM report (Technical Panel on Assumptions and Methods, 2003). My thanks go to Alice Wade of the Social Security Advisory Board for running this simulation.

Table 11. Actual and Predicted Growth Rates of Components of Real GDP and Related Variables

Percent

Component or related variable	Actual, 1987:3–2001:2		Projected, 2003–23	
	Total economy	Nonfarm private business sector	Total economy	Nonfarm private business sector
Components				
Output per hour	1.28	1.78	2.00	2.50
Hours per payroll employee[a]	–0.15	–0.15	–0.10	–0.10
Employment rate	0.12	. . .	0.00	. . .
Labor force participation rate	0.13	. . .	0.00	. . .
Working-age population	1.16	. . .	0.98	. . .
Mix effect	–0.51	. . .	–0.20	. . .
Payroll-household employment ratio	0.47	. . .	0.10	. . .
Related variables				
Aggregate hours	1.36	1.46	0.98	1.08
Payroll employment	1.82	1.61	1.08	1.18
Household employment[b]	1.41	. . .	0.98	. . .
Hours per household employee	0.26	. . .	0.00	. . .
Output	2.95	3.24	3.28	3.58

Source: See appendix A.

a. It is assumed that aggregate hours per payroll employee changes at the same rate in the total economy as in the nonfarm private business sector as measured by the Bureau of Labor Statistics.

b. Growth in payroll employment minus growth in the payroll-to-household employment ratio.

causing an increase in the growth rate of payroll employment relative to household employment. But more of these jobs are part-time, causing a reduction in hours per payroll employee. These two effects did not quite cancel out, in that hours per household employee rose slightly during 1987–2001. The 2003 TPAM report interprets this increase as temporary and predicts a growth rate of hours per household employee of exactly zero over the next seventy-five years. This paper adopts their assumption.

Implications for Potential Real GDP Growth

Table 11 provides a template for forecasts of potential GDP over the next twenty years. I base the projections on the analysis above, plus the growth rates of the components of the output identity over the period between two benchmark quarters, 1987:3 and 2001:2. Here the two left-

hand columns display for this period the growth rates of real GDP and the components of the output identity, plus growth rates of related variables that can be extracted from the ratios in the identity or from published information.

Forecasting potential real GDP growth for 2003–23 requires a projection for each component of the output identity. I have already settled on a future growth rate of NFPB output per hour of 2.50 percent a year. Also, the assumption of zero growth in hours per household employee implies that the growth of hours per payroll employee must be exactly offset by the growth of payroll employment relative to household employment. I choose growth rates of –0.10 and +0.10 percent a year for these two off-setting factors. Offsetting the projected decline in the growth rate of hours per household employee from +0.26 percent a year in 1987–2001 to zero during 2003–23 is a projected decline in the growth rate of the mix effect from –0.51 percent in 1987–2001 to –0.20 percent in 2003–23. I see no reason to forecast changes in the employment rate from its value of roughly 95 percent in early 2001, or in the labor force participation rate from its early-2001 value. Numerous offsetting trends influence the LFPR. The movement of females into the labor force, which boosted the LFPR between the 1960s and 1980s, appears to be largely completed. The long-standing shift to earlier retirement ages, especially for adult men, also appears to have stabilized or even reversed.[66] This paper's estimated trend for the LFPR registers a growth rate of zero between 1995 and 2003 (table 9). Thus I project a zero growth rate for both the employment rate and the LFPR.

The earlier discussion of population growth emerged with a projected 2003–23 growth rate of 0.98 percent a year. This completes the list of components to be forecast in order to project potential real GDP; the predicted growth rate that emerges from these component forecasts is 3.28 percent a year over the next two decades. Even though NFPB productivity is forecast to grow 0.72 percentage point faster than during 1987–2001, potential real GDP is forecast to grow only 0.33 percentage point faster. The difference is explained by a shift from an increasing employment rate (a declining NAIRU) to a stable employment rate, a shift from a slight increase in the LFPR to a stable LFPR, and a decline in population growth by 0.18 percentage point a year.

---

66. See Technical Panel on Assumptions and Methods (2003, figure 11, p. 52).

Conclusion

Productivity growth experienced a second acceleration in 2000–03 following the initial productivity revival of 1995–2000. This recent upsurge has brought good news for the American standard of living, bad news for many American workers made jobless by a continuing erosion of payroll employment, and new puzzles for academic research. This paper has placed the 2000–03 productivity acceleration at center stage and illuminated it from the three perspectives of postwar history, possible causes, and prospects for future economic growth.

The paper began with three goals: to forecast future potential GDP growth for two decades into the future, to explain the recent productivity upsurge, and to assess the recent productivity performance in the context of trends and cycles over fifty years of postwar history. Both future projections and historical analysis require the linking of potential GDP growth to productivity growth through an output identity that identifies seven variables, including productivity and the employment rate; the product of these seven variables must by definition equal real GDP. The paper's analysis began with history, developing an econometric decomposition of all seven components of the output identity into trends, regular cyclical responses, and residual unexplained responses. It then proceeded to examine possible explanations for the recent productivity upsurge, including both substantive hypotheses and possible sources of measurement error. The paper's last section projected components of the output identity out over two decades, yielding a forecast of future potential real GDP growth.

As the first step in the empirical analysis, the paper established trends for real GDP and the components of the output identity over the period since late 1954. It set the trend of the employment rate as a log-linear spline trend between benchmark quarters having similar characteristics. For productivity growth and the other five components of the output identity, the paper used an average of results from a Hodrick-Prescott filter and a Kalman filter that estimates a time-varying coefficient in an equation that controls for output fluctuations. The implied trend for potential real GDP was obtained by multiplying together the trends for the seven components of the output identity. The most important conclusion was that the estimated trend of productivity growth in 2002–03 was slightly above 3.0 percent a year, substantially above the recent consensus of pro-

ductivity experts and indeed higher than the previous maximum trend rate achieved in the early 1960s. The corresponding trend in potential real GDP in mid-2003 was 3.38 percent a year.

Armed with these trends, the paper developed a dynamic econometric analysis in which changes in deviations from trend for each component of the output identity depend on current and lagged changes in the output deviation from trend, a set of lagged dependent variables, an error correction term, and an end-of-expansion effect. Each component exhibited a highly significant positive or negative response to output fluctuations, and the end-of-expansion coefficients told a consistent story in which labor input grows too rapidly relative to output in the final stages of an expansion, while productivity grows too slowly, and in which this overhiring is eliminated in the recession and early stages of the recovery. These results yielded another important conclusion: that the econometric coefficients based on nearly five decades of postwar history fail to explain why actual productivity growth exceeded trend growth in 2002–03; a large unexplained residual emerges as the difference between rapid actual growth and a prediction that productivity growth should have been below trend due to the sluggish below-trend rate of output growth during this interval. Since the same regression is estimated for all seven components of the output identity, one can trace the counterpart of positive productivity residuals to negative residuals for other components of the output identity.

In previous business cycles, the first four recovery quarters shared in common with 2002–03 a very rapid rate of productivity growth relative to trend, which I have called the "early-recovery productivity bubble." In every past episode the large positive excess of actual over trend productivity growth was followed by a return to trend or even below-trend growth in the following eight quarters. Although the econometric equation can explain the earlier bubble episodes with small statistical residuals, the positive residual is very large in 2002–03, again confirming that the recent upsurge was unusual by every measure, including the acceleration of the estimated trend, the positive deviation of actual growth from trend, and the inability of the econometric coefficients to capture the dynamics of changes in the deviation of actual from trend.

If the recent productivity upsurge was outside the bounds of normal cyclical behavior, what explains it? Two complementary hypotheses occupy center stage. First, the upsurge was preceded by an unusual trajectory of corporate profits, which grew rapidly in 1998–2000 and then

collapsed in 2001 as the result of shady accounting, the revelation of accounting scandals, and massive write-offs. The extraordinary profit squeeze led in turn to savage cost cutting and layoffs, reducing labor input long after real GDP had begun to recover in late 2001. But, second, firms could not have maintained or increased output following the reduction in labor input without the help of both measured ICT investment and unmeasured complementary investments in intangible capital. Recent research shows that intangible capital that is complementary to computer hardware and software can cause measurement errors in official data on productivity, leading productivity growth to be understated in a period of buoyant ICT investment such as 1995–2000, and to be overstated in a period of slumping ICT investment such as 2000–02. Both the profit squeeze hypothesis and the intangible mismeasurement hypothesis imply that the supercharged productivity growth rates observed in 2002–03 are inherently temporary and will soon return to the still-respectable estimated trend of around 3.0 percent.

Projections of potential GDP growth for the next two decades require forecasts of growth in productivity, population, hours per employee, and the other components of the output identity. I have argued that the underlying trend in productivity growth will decline from the recent rate of 3.0 percent a year to about 2.5 percent in the next two decades, because of the unwinding of temporary factors bolstering recent productivity growth, and because of potential diminishing returns in the exploitation of New Economy innovation, as well as the implications of a plateau in educational attainment for sharply slowing future growth in labor quality. Population growth is forecast to slow from 1.2 percent a year over the past decade to around 1.0 percent, which is still a higher forecast than that embedded in current Social Security projections. Hours per employee (as defined by the household survey), the employment rate, and the labor force participation rate are all forecast to be unchanged over the next two decades.

Taken together, these assumptions yield a projection for potential real GDP growth of 3.28 percent a year over 2003–23. This is much slower than the rate forecast by the optimists cited at the beginning of this paper, who thought (or still think) that 4.0 percent annual growth of potential output is feasible, but more optimistic than the future growth assumptions embodied in the annual report of the trustees of the Social Security Administration. Along the way, this examination of future trends has

required extending the analysis from the traditional turf of productivity forecasting to the less familiar terrain of fertility rates, mortality rates, and immigration assumptions.

In the end, this exercise in crystal ball gazing is a humbling experience. Although forecasts of the employment rate or the labor force participation rate over the next two decades may have only a small margin of error, the degree of uncertainty is much larger for three key components of the output identity, namely, growth in productivity, in population, and in hours per employee. Neither economics nor statistics provides much guidance about the measurement of uncertainty in making these forecasts, and part of this uncertainty concerns how much of the past to use in making future projections. Other demographers cited here have, within the same study, taken 180 years of history as relevant for immigration, 50 years of history as relevant for mortality, but only 20 years of history as relevant for fertility. Is it surprising, then, that economists disagree on whether the next two decades of productivity growth is better represented by average growth over the past two years, or the past eight, or the past twenty? Perhaps the only safe forecast is that we will be surprised sooner rather than later, just as the recorded productivity growth rates of the past two years would astonish and bewilder any Rip van Winkle who had fallen asleep in the autumn of 2001.

APPENDIX A

Data Issues

Data Sources

All data were obtained from the websites of the Bureau of Labor Statistics (BLS; www.bls.gov) and the Bureau of Economic Analysis (BEA; www.bea.gov), were current as of September 15, 2003, and were retrieved for the period 1948:1 to 2003:2. Data subsequently released for 2003:3 and any revisions subsequent to September 15, 2003, are not included in the analysis.

The BLS data are identified below by their symbol in the paper (see equation 2) and by their BLS series identifier. "NFPB" refers to the nonfarm private business sector:

NFPB employment (E^B)   PRS85006013
NFPB aggregate hours of work (A^B)   PRS85006033
NFPB output (Q^B)   PRS85006043
Nonagricultural payroll employment   CES0000000001
Civilian noninstitutional population (N)   LNU00000000
Civilian labor force (L)   LNS11000000
Civilian employment (E)   LNS12000000

NFPB output per hour was calculated as Q^B/A^B, and NFPB hours per employee was calculated as A^B/E^B.

Initial research for this paper used private nonagricultural payroll employment from the Current Employment Survey (CES0500000001) as a proxy for NFPB employment. The preferred measure for E^B listed above is not listed in the quarterly news release of the BLS productivity data but is available on the BLS website. It differs from the CES measure in several ways: it includes part of private nonfarm employment not covered by the CES, as well as employees of government enterprises, nonfarm self-employed workers, and nonfarm unpaid family workers, and it excludes employment in nonprofit enterprises. The two sources grew between 2000:2 and 2003:2 at annual rates of –0.72 percent for the CES series and –1.04 percent for the preferred NFPB series, implying an overstatement of the decline in NFPB hours per employee when the CES series is used.

The BEA website was the source for quarterly real GDP (table 1.2) and annual agricultural employment (table 6.8c), which was interpolated from annual to quarterly. Data used for total payroll employment E^P are the sum of this interpolated quarterly series on agricultural employment and the official BLS series listed above for nonagricultural payroll employment.

Smoothing of Jumps in Household Data

BLS data on population, civilian labor force, and civilian household employment exhibit periodic jumps when they are adjusted for the latest census results. To eliminate the effect of these jumps, I calculated a smoothing adjustment for six quarters when the annual growth rate of the population exceeded 2.5 percent and this rate of change appeared to be a clear outlier from neighboring data. These outlier quarters were 1953:1, 1960:1, 1972:1, 1990:1, 2000:1, and 2003:1. The smoothing exercise was

carried out for the population data, and the same percentage adjustment factors were then applied to the labor force and employment data. Thus the smoothing affected all data reported in the paper for N and E^p/E, whereas the ratios L/N and E/L were not affected. Excess population growth in the outlier quarters was calculated by subtracting the average quarterly growth rate during the four quarters preceding the jump quarter. The excess growth, in all cases except 1953:1, was then spread evenly over the 40 quarters preceding the jump by multiplying the original population in those quarters by e^{nx}, where n indexes quarters, starting at 0 in the 40th quarter before the target quarter, and x is the excess growth in the target quarter divided by 40. For 1953:1 the excess was spread over the 39 following quarters by multiplying them all by $e^{(40-m)(-x)}$, where m indexes quarters, starting at 0 in the target quarter, and x is the excess growth.

Definition of Benchmark Quarters

The trend-through-benchmarks (TTB) detrending method requires the choice of benchmark quarters to divide the periods. To establish quarters that represent roughly the same degree of cyclical pressure in each business cycle, I chose those when the actual unemployment rate was equal to the NAIRU, where the latter is estimated in Gordon (1998), as updated in Eller and Gordon (2003). In each business cycle, starting at the trough, the unemployment rate declines through the NAIRU and in the next recession rises through the NAIRU; the benchmark quarters occur at the first of these events, the quarter when the actual unemployment rate first declines below the NAIRU.

Table A1 lists the benchmark quarters. For convenience I have taken the pre-1993 benchmark quarters from the list in Gordon (1993, table 3) and have added 1994:4. As yet there has been no quarter during the current recovery when the unemployment rate has fallen below the NAIRU, and so I take 2001:2, a quarter during the most recent recession when unemployment was rising through the NAIRU, as the terminal point for calculating growth rates. The choice of 1994:4 is easy: it is a period like 1987:3, when the economy was expanding at a moderate pace with subdued inflation. The choice of the terminal date 2001:2 is more problematic. Because the unemployment rate had fallen so low in 1999–2000, it had not yet risen as far as the NAIRU even in 2001:2, in the middle of the

Table A1. Unemployment Rate in Benchmark Quarters, 1954–2001

Business cycle[a]	Benchmark quarter[b]	Unemployment rate (percent)
1953–57	1954:4	5.4
1957–60	Excluded	…
1960–69	1963:3	5.5
1969–73	1972:2	5.7
1973–80	1978:3	6.0
1980–81	Excluded	…
1981–90	1987:3	6.0
1990–2001	1994:4	5.6
Terminal date	2001:2	4.8

Sources: Bureau of Labor Statistics data; Gordon (1993, p. 285).
a. Measured from peak to peak.
b. Selected as described in the text. The short business cycles of 1957–60 and 1980–81 are omitted because actual unemployment fell below the NAIRU in 1959–60 only briefly and by a trivial amount, and in 1980–81 actual unemployment never fell below the NAIRU.

recession. (Eller and Gordon, 2003, estimate the NAIRU to have been in the range of 5.2 to 5.4 percent in 2001.) Yet to place the terminal date later would yield readings on labor force participation and hours per employee that reflect weak labor market conditions and are not comparable to the earlier benchmark quarters.

Timing of End-of-Expansion Dummy Variables

The end-of-expansion (EOE) dummy variables are shown in table A2, where k represents the number of the dummy (corresponding to the coefficients shown in tables 4 and 5). As explained in the text, each of the seven dummy variables equals $1/M$ during the "on" quarters shown in the

Table A2. Timing of End-of-Expansion Dummy Variables, 1955–2003[a]

Dummy variable (k)	M	D(k) = 1/M during	N	D(k) = –1/N during
1	7	1955:4–1957:2	4	1957:3–1958:2
2	8	1958:3–1960:2	8	1961:1–1962:4
3	7	1968:2–1969:4	6	1970:2–1971:3
4	7	1973:1–1974:3	7	1974:4–1976:2
5	11	1978:4–1981:2	4	1982:3–1983:2
6	8	1988:2–1990:1	8	1991:1–1992:4
7	5	2000:1–2001:1	7	2001:4–2003:2

a. The $D(k)$ variables would be entered as $1/M$ and $-1/N$ if the data were annual. With quarterly data, as in this paper, they are instead defined as $4/M$ and $-4/M$.

table and $-1/N$ during the "off" quarters, where M is the length in quarters of the initial interval of excessive labor input growth, and N the length of the subsequent correction. With a few minor exceptions, the timing of the first five variables corresponds to that listed in Gordon (1993, p. 291), and the sixth to the "early EOE effect" discussed in Gordon (1993, p. 300). The "on" period for 2000–01 was chosen to begin after the peak quarter of real GDP growth (1999:4) and to extend to the peak quarter as identified by the National Bureau of Economic Research (2001:1). The "off" period was chosen to begin in the first quarter of positive real GDP growth in the recovery (2001:4) and extend through the end of the sample period.

Comments and Discussion

Martin Neil Baily: This paper continues the tradition of Robert Gordon's insightful evaluations of productivity and potential GDP growth, two related issues that are among the most important in empirical macro-economics. Gordon devotes most of the paper to figuring how to separate out the trend rate of productivity growth from the cycle. This is a tough assignment given what appear to be rather sharp shifts in the trend in the early 1970s and in the mid-1990s, overlaid with cyclical effects that look very different in the past two cycles. Gordon's views about productivity have changed over time, and appropriately so, given the new data and the revisions of older data that have emerged.

As part of the process of separating trend from cycle, Gordon reexamines the findings of the classic paper by Arthur Okun that introduced the concept of potential output and first examined a puzzle in its behavior: why fluctuations of actual around potential GDP are so much larger (by a 3-to-1 ratio in Okun) than those of the unemployment rate around its baseline, what we now call the NAIRU, or natural rate of unemployment. Finally, Gordon uses his assessment of the productivity growth trend to examine what assumption should be made about potential GDP growth into the future. This is a key issue for the future of Social Security.

The paper explores three approaches to identifying trend and cycle. The first uses benchmark periods, for which it is judged that the business cycle was at about the same point at the beginning and the end of the period. Alternatively, a regression analysis can be done in which right-hand-side variables pick up the cyclical variations in the time series, allowing for an adjusted or trend value to be separated out. Finally, filters can be used to smooth a cyclical series. Both the Kalman filter and the

Hodrick-Prescott filter are widely used in the literature. The latter is a flexible approach that allows the researcher to select from a range of smoothing parameters, depending on the series to be smoothed. That is both a plus and a minus, of course, since one sometimes gets a sense that the parameters are being chosen on the basis of some prior belief about the right answer.

Picking benchmark dates has the danger that it becomes an endogenous process. The researcher may find it difficult to select dates without peeking at the impact that different dates may have on the results he or she is looking for. One solution is to use the dates selected by the Business Cycle Dating Committee at the National Bureau of Economic Research, although the committee looks at the same output and employment data that go into measures of GDP and productivity. The regression approach has appeal if one could find both a sensible specification and right-hand-side variables that are exogenous. I am probably more willing than some econometricians to accept that a particular variable is exogenous, but in the case of aggregate productivity regressions it is very hard to find variables that even come close to meeting a plausibility test. Most candidate variables are related very closely either to output, which is the numerator of any labor productivity measure, or to employment, which is related to the denominator. In his regression approach, Gordon has used a variant of output as his right-hand-side variable, and his regressions have a corresponding tendency to overestimate the cyclical component of productivity variations, as Gordon himself notes. It would be helpful to contrast Gordon's findings with those of the Council of Economic Advisers (CEA), which has used a labor demand equation instead of a productivity regression, and the study by Basu, Fernald, and Shapiro, cited by Gordon, which uses yet a different approach.[1] Those two studies found that the post-1995 acceleration in productivity was not at all a cyclical phenomenon.

Filtering approaches have been used in work at the Federal Reserve Board and in the big cross-country study of growth at the Organization for Economic Cooperation and Development (OECD). My concern about them is that they impose a particular pattern on the productivity data that does not appear in the basic numbers. Since the filters avoid slope discontinuities, they look ahead at an upcoming break in the trend and assume that it must have started before the apparent date of the break point. I see

1. Basu, Fernald, and Shapiro (2001).

nothing in the actual data for the 1990s to suggest that productivity began to accelerate before 1996, but in Gordon's results using filtering methods, and indeed in the OECD's similar estimates for the United States, the productivity acceleration started in the early 1990s. Maybe that is correct. I have to admit that there seems to be no smoking-gun explanation of the shift in the trend that can account for a sudden trend break in 1996. But it would be more convincing if we could understand what forces kept productivity growth below trend in the early 1990s.

A piecewise linear time trend fits the labor productivity data in the United States remarkably well. In my own work, I have simply imposed break points on the data, but Steven N. Braun in his work at the CEA has found it very hard to do as well as a piecewise linear time trend in fitting these data. A recent paper by James Kahn and Robert Rich has explored this possibility further by using growth theory to suggest additional variables, including real wage movements, to help separate trend and cycle.[2] These variables are measured using different source data and will have different random errors. The Kahn-Rich fitted trend productivity growth rate comes very close to a piecewise trend.

Regardless of which adjustment approach one uses, it does now seem clear that trend labor productivity growth picked up substantially in the 1990s, and the most recent data suggest that there may even have been a further acceleration in the past two years. In the data available as this goes to press, output per hour in nonfarm business grew at an astounding 4.4 percent annual rate from 2001:1 to 2003:3. The rate for manufacturing is similar, indicating that rapid growth is continuing in the nonfarm, non-manufacturing sector. Like Gordon, I am wary of the most recent data. In principle, revisions up and down should be equally likely, but in practice the GDP and productivity numbers were revised down in 1999, 2000, and 2001. This may indicate that some systematic error has entered the estimation procedure and is yielding too optimistic a view of the economy in preliminary numbers.

One test of the 2002 data will come in December 2003 when the GDP accounts are revised. We will find out more about 2003 with the revisions due in the summer of 2004. But unless the data are changed quite drastically, it looks as if the productivity acceleration has survived its first business downturn and, moreover, has survived the slump in technology

---

2. Kahn and Rich (2003).

spending. That is a big deal. It suggests that the acceleration was not just a result of strong demand or of the wave of spending on information technology (IT). Gordon accepts the first point but disputes the second.

On the question of whether IT is the reason for the productivity acceleration, I have been on both sides at different times. In 1988 I co-wrote a Brookings Paper with Gordon on the role of IT and its interaction with measurement issues and gave a variety of reasons why IT would not show up in measured productivity growth. Measurement methods have improved, but many of the same arguments made then still hold today.

In the late 1990s it seemed obvious that IT was the most important factor driving the productivity acceleration. All the evidence seemed to point the same way, including sharp increases in the number of patents and in the volume of R&D as well as the surge in IT investment. A range of studies using growth accounting techniques and other approaches supported that view. Since 2000 I have scaled back my estimate of the importance of IT to the productivity acceleration for the following reasons. First, growth accounting decompositions, however valuable, do not prove that the hedonically inflated IT capital is really being used productively. They simply assume that is the case. Some reverse causality may also have been at work in which the 1990s boom (in part) caused the surge in IT spending rather than vice versa. Second, industry case studies at the McKinsey Global Institute have found that there were many important ingredients besides IT that caused productivity to rise in the late 1990s.[3] Third, industry analysis by Jack Triplett and Barry Bosworth concluded that the service industries that had experienced the largest increases in labor productivity were the ones that had experienced the largest increases in TFP.[4] And fourth, the continuation of very rapid growth after 2000 has broken the simple correlation that emerged in the late 1990s, where faster productivity growth coincided with rapid growth of IT capital.

On balance, IT is indeed an important enabling technology that, when accompanied by business process innovation, can allow companies to raise their productivity. But the important thing is for companies to figure out how to improve their products or processes. Those that buy IT capital without a clear understanding of how it will improve operations tend to find little or no payoff.

3. McKinsey Global Institute (2001).
4. Triplett and Bosworth (2002).

One reason for the rapid productivity growth of the 1990s that continued to operate strongly through the 2001 downturn is that companies faced very strong competitive pressure with little effective resistance from labor. An interesting finding in Gordon's paper is that the cyclical pattern of productivity looks different since the mid-1980s. The shifts in both trend growth and cyclical responsiveness indicate that there may have been structural changes in the economy brought about by a combination of deregulation, globalization, and a loss of labor's bargaining power, with the following consequences. They have made it both easier and more essential for companies to make changes in their operations, including outsourcing. They have forced companies to lay off workers who are not needed immediately. And they have placed very strong performance pressures on individual employees.

In his discussion of intangible capital and the likely delays in achieving the potential impact of IT capital, Gordon addresses some of these same issues. He makes a valuable contribution toward extending the growth accounting framework and making it more realistic. There is still a way to go before we have a complete formulation, but he provides an important start on the problem. One must be careful about interpretation, however. Altering the growth accounting framework may make its results more consistent with some facts—such as the persistence of rapid productivity growth even as technology spending slumped. But that comes at a price. It was the coincident timing of heavy IT spending and productivity acceleration from 1995 to 2000 that made so many people into converts to the New Economy. That evidence is weakened once one allows for delayed effects and the need for intangible capital.

Turning to the issue of the long-run growth estimates used by the Social Security Administration and the difficulty of estimating how fast GDP will grow over the long run, I first want to point to a striking regularity. Over the twentieth century as a whole (1899 to 1999, to be precise), GNP per capita in the United States grew at 2.1 percent a year on average. Over the first half of the century, growth averaged 2.0 percent a year, and over the second half it was 2.2 percent. There was some variation over subperiods, notably during the Great Depression. But there was less variation than one might imagine during the second half of the century. GNP per capita grew at an average rate of 2.4 percent a year over 1948–73, at 1.9 percent over 1973–90, and at 2.1 percent over 1990–99. These figures all are taken from the *Economic Report of the President, 2000,* and my

colleague on the CEA at the time, Robert Lawrence, dubbed this finding the "2 percent solution." If you want to know how fast the U.S. economy will grow over the long run, the answer is 2 percent per capita per year, give or take a little. Of course, as the prospectuses always say, the past may not be a guide to the future, so it is helpful also to look, as Gordon does, at the ingredients that go into output growth: population, labor force, hours worked per employee, and productivity.

Gordon attributes much of the recent growth to IT, which suggests that, once Moore's Law has run its course, productivity growth could slow sharply. However, he remains somewhat optimistic, suggesting that 2½ percent a year is an appropriate forecast over the next twenty years for labor productivity growth in the nonfarm business sector. That is a pretty good number. My own estimate of the current productivity growth trend is that it is running at around 2½ percent a year, or maybe 2¾ percent if the recent data hold up. But if forced to forecast twenty-five years into the future, I would be a bit more cautious than Gordon. It is certainly possible that the economy will return to something closer to the 1.4 percent annual growth rate that the United States experienced for over twenty years. Alternatively, if one takes the very long term as a guide, business sector output per hour grew at 2.1 percent a year on average over the twentieth century, and the farm sector contributed positively to that growth. So non-farm business productivity was probably around 2 percent a year. I would be comfortable using a projection of 2½ to 2¾ percent a year for the next five years. For a seventy-five-year projection, I would use 2 percent a year. For a twenty-year horizon, 2¼ percent growth looks sensible.

On population and immigration, there are more questions than answers. The impact of immigration will depend on the productivity of the people who come and how rapidly they assimilate into the American economy. On the first of these, immigration into the United States is bimodal: many individuals of high ability come to work in programming and other professions, but at the same time there appears to be a whole class of low-wage jobs that Americans do not want and are now filled by immigrants. Of course, the same was true in earlier periods of high immigration. And it may be that the next generation of immigrants will acquire skills and education and move up the wage and productivity ladder, but that is an open question.

After redoing Okun's analysis, Gordon concludes that, once lags are considered, the Okun coefficient is closer to 2 than to 3. That result seems

correct, if not surprising given estimates made since Okun's original study. Gordon finds the NAIRU to be between 5.2 and 5.4 percent, very close to the figure of 5+ percent that has prevailed in CEA forecasts spanning several administrations. Whether the NAIRU is a meaningful concept is still debated. For me the idea that at any given time there is a constraint on the ability of demand policies to increase output without also increasing inflation seems highly plausible and is consistent with Okun's original analysis of potential output. Presumably, at any given time there is an unemployment rate associated with the inflation constraint, but there is no guarantee that the NAIRU defined in this way will remain fixed over time. One criticism of the NAIRU, however, is surely correct. Even if the unemployment rate were to stay stuck above the NAIRU, it does not seem remotely plausible that inflation would decline at an ever-faster rate.

Finally, there is the issue of American exceptionalism. Is it possible that the United States will continue to have faster productivity growth than Europe for an extended period? That is an interesting question, and one that is related to a whole branch of research in recent years. Gordon implicitly has a convergence model in mind: Technologies developed in one country can be used in others. If capital is in short supply, it will flow in from outside. But we also know that convergence is not guaranteed. Many—indeed, most—of the countries of the world have not converged toward U.S. levels of GDP per capita or productivity. Technologies and business practices used in one country do not necessarily flow to all countries.

An emerging view in the literature suggests that institutions, legal arrangements, property rights, market flexibility, regulation, and economic policies are key factors in determining whether convergence takes place. Sometimes this is posed as a debate between the institutional view and the policy view, but the two are not orthogonal. The paper by Barry Bosworth and Susan Collins in this volume is relevant to that discussion.

On Europe specifically, although it is important to remember that not all European countries are alike, the larger continental countries have gone down a different economic road than the United States. Their economies have less flexibility, less competition, and a more restrictive set of labor market rules. The result has been a society with less economic insecurity than the United States, but also less employment and somewhat

lower productivity in the private business sector.[5] The extent to which Europe and the United States will converge economically depends on whether Europe is willing to undertake reforms.[6]

This was an enjoyable paper whose conclusions have been developed and sharpened as it has progressed. It did not raise my blood pressure, which has not always been true of Gordon's papers in the past. That is good in that it suggests broader consensuses are emerging, even though it makes for less combative comments.

Daniel E. Sichel:[1] Continuing a long-standing tradition of the Brookings Panel, Robert Gordon has written another interesting paper on productivity, and I am happy to have the opportunity to discuss it. Gordon does three big things in this paper. First, he sets up the usual decomposition of output growth into its supply-side components, detrends each piece of the decomposition, and estimates a dynamic error correction model for each. Second, he uses this machinery to examine several interesting questions: How unusual is the pattern of labor productivity growth since 2000? What is the current trend in that growth? How has the Okun's Law relationship held up recently? Given his estimates that labor productivity growth in recent years was much stronger than predicted by his dynamic model, Gordon explores possible reasons for the unusual performance. Finally, he puts together a plausible projection of potential output growth for the next twenty years.

This paper covers a lot of ground, and Gordon provides some useful insights into the recent behavior of productivity. In my judgment, the paper proposes some plausible explanations for the surprising strength in productivity growth over the past couple of years, but it does not nail down the likely quantitative contribution of these alternative explanations. Here I will focus primarily on Gordon's estimates of trend productivity growth

5. Based on GDP per hour calculations, France and Germany have labor productivity as high as or higher than the United States, but estimates of business sector productivity show a productivity gap in the United States' favor. See, for example, O'Mahony and de Boer (2002).

6. See my book on European reform, to be published in the spring of 2004 (Baily, forthcoming).

1. The views expressed in this comment are the author's alone and should not be attributed to any member of the Federal Reserve Board or its staff.

and his discussion of the recent strength in productivity. I will offer a few comments on the other elements of his paper as well.

I will start with Gordon's estimate of trend productivity. He uses an average of a trend from an H-P filter with a smoothing parameter of 6,400 and a trend estimated with a Kalman filter model. This procedure suggests that trend productivity has increased at an annual rate of about 3 percent from 2001 to the present. Given his past pessimism on productivity growth, an estimate of this magnitude from Gordon is noteworthy. It is considerably faster than many other researchers have suggested recently, and it points to a dramatic shift in the productivity debate. As recently as mid-2002, the debate still centered on how much of the post-1995 pickup was permanent and how much was transitory.[2] Now the debate is focused on how much further trend productivity growth might have increased beyond its 1995–2000 pace.

Although both the H-P and the Kalman filters have been used extensively and may be delivering a sensible estimate, some issues can be raised with Gordon's procedure. The H-P filter suffers from well-known endpoint problems, which make its uses somewhat problematic for obtaining trend estimates for recent years. Also, it is a univariate procedure and as such does not take account of information in other economic variables. In addition, the choice of the smoothing parameter is arbitrary. The simple type of Kalman filter model that Gordon estimates potentially avoids these problems: it is possible to estimate the smoothing parameter and to introduce information from other variables. However, one issue may remain. Because the latest observations on output—which are crucially important to the estimate of current trend productivity—have not been through an annual revision, they are likely to be more subject to measurement error than earlier observations. If so, it might be appropriate to reduce the influence of the most recent observations on estimates of trend productivity. Given these technical issues, it might be better to avoid the H-P filter and refine the estimates from the Kalman filter alone.[3]

Conditional on this estimate of trend, Gordon uses his dynamic model to judge the size of the surprise in productivity growth relative to his esti-

2. As an example of the tenor of the debate in mid-2002, in August of that year Goldman Sachs published a piece that began as follows: "Rest in peace, New Economy: it was fun while it lasted" (Goldman Sachs, "RIP, New Economy," *U.S. Economics* newsletter, August 14, 2002).

3. For a more full-blown Kalman filter model of trend productivity, see Roberts (2001).

mated trend. But he does not tell us about the surprise in the trend itself. In any case, the results in Gordon's table 8 suggest that productivity has increased about 1½ percentage points a year faster since 2001:4 than the model can explain.[4] However, if one believed that the current trend in productivity growth were less than 3 percent, then the surprise would be larger; conversely, if one believed that the current trend were greater than 3 percent, the surprise would be smaller.

Gordon proposes several explanations of the surprise in productivity growth. Of the two he finds most compelling, the first is "that an unusual degree of downward pressure on profits led to unusually aggressive cost cutting by firms." I find it plausible that firms have squeezed extra efficiency out of their current employees, although they may have done this for reasons other than the profit-squeeze story. (For example, greater-than-usual uncertainty about the timing and strength of the recovery may have led to greater-than-usual reluctance of firms to hire. Or, less neoclassically, firms became fat and lazy during the boom years and have been squeezing out those inefficiencies in the past couple of years.) It is important to emphasize that explanations of this sort suggest, at best, a one-time shift in the *level* of productivity. That is, once these efficiency gains have been achieved, there is no reason to believe that further gains of this sort are possible or that trend productivity growth will get a boost from this source in the longer run. Indeed, it is even likely that a portion of the gains that come from squeezing workers harder will be reversed once the labor market tightens. In any case, Gordon does not quantify the productivity gains that might be coming from firms achieving these sorts of one-time increases in efficiency.

Gordon's second leading explanation hinges on two elements: a delay between the acquisition of information technology (IT) capital and its productivity effect (in the spirit of Paul David's analogy to the delayed productivity impact of electric motors a century ago), and a surge in purchases of unmeasured intangible capital in the late 1990s. The basic idea is that, in the late 1990s, firms invested heavily in IT-related intangible capital—which was not counted in the official output statistics—and that that capital recently became productive and has boosted measured productivity growth in the past couple of years. Gordon suggests that,

4. I obtain this estimate by taking a weighted average of the residual over 2001:4–2002:3 (2.02 percentage points) and the residual over 2002:4–2003:2 (0.75 percentage point).

because of this effect, labor productivity growth could have been under-stated by 0.5 to 0.6 percentage point a year in the late 1990s and that it is being overstated by somewhat less now.

This calculation starts with an assumed investment share for these intangible investments and an assumption that they grew at the same rate as *real* investment in IT during 1995–2000. (For example, the paper notes that the growth rate of real investment in computer hardware and software was 22.5 percent a year over this period.) It seems plausible to me that investment in IT-related intangibles could have increased at the rate of growth of *nominal* IT investment; for computer hardware and software together, that rate averaged about 13 percent from 1995 to 2000, accord-ing to estimates by the Bureau of Economic Analysis.[5] But it is harder for me to believe that prices for intangibles are falling anywhere near as fast as prices for tangible IT; that is, it is hard to see how organizational capi-tal could undergo as rapid a pace of technical change as, say, computer hardware. Thus it appears likely that *real* spending for intangibles rose more slowly in the late 1990s than real spending for IT. If so, the esti-mated effect of the intangibles-and-delay hypothesis on productivity growth in recent years would be smaller than suggested by Gordon's estimates.

More broadly, Gordon's story about intangible capital and delay seems qualitatively plausible to me. Quantitatively, I am at more of a loss. For the reason described above, I am unsure that Gordon's estimates are right. Moreover, I would encourage further specificity about the types of intan-gible investment that have been delayed and the specific sources of that delay. Such specificity, along with an actual model, would make it easier to assess the empirical magnitude of the various delay hypotheses.

Gordon also considers, and largely dismisses, other possible explana-tions of the productivity surprise. On the measurement front, these include the argument that the Bureau of Labor Statistics' household sur-vey gives a more accurate picture of employment during a cyclical recov-ery than does the Bureau's payroll survey, as well as the story that unmeasured hours of work have increased substantially recently. For the

---

5. As another reference point, a recent paper by Corrado, Hulten, and Sichel (2003) estimated that nominal spending on intangibles excluding software (which is already counted in official GDP measures) increased at an annual rate of 8½ percent from the mid-1990s to the late 1990s, very similar to the rate of increase in nominal outlays for computer hardware.

reasons that Gordon lays out, I am inclined to agree that these are unlikely to be important parts of the explanation of the recent productivity surprise. On the household versus payroll debate, I would add that the Bureau recently changed the procedures for the payroll survey in a way that should improve its ability to track employment during cyclical recoveries. One other measurement issue that should be mentioned is that the analysis in the paper is based on GDP data from before the benchmark revision of the National Income and Product Accounts in December 2003. It is always possible that this revision will change the productivity picture.

Gordon also explores the possibility that greater outsourcing of white-collar jobs has been a source of productivity improvement in recent years. I agree with Gordon's conclusion that this factor is unlikely to explain a large part of the pickup in productivity growth in the past couple of years or in the mid-1990s. However, I believe that the channel through which outsourcing affects productivity growth could be more complicated than the example in the paper suggests. Indeed, it is even possible (although perhaps not likely) that outsourcing could lower U.S. productivity growth as well as raise it.

For example, consider the simple case in which high-productivity workers (and the capital they work with) are outsourced. In this case, all else equal, the average productivity level of workers still in the United States would be lower than before, because there would be fewer high-productivity workers in the United States. Much of the recent media coverage of outsourcing has focused on white-collar jobs, some of which are likely to be high-productivity jobs such as those of software engineers. Of course, many factors would affect the impact of outsourcing on productivity, including what happens to the capital that the outsourced worker had been using. If that capital stayed in the United States to be used by other workers, the additional capital per worker would tend to increase the productivity of the remaining workers. If instead that capital were outsourced along with the labor, the extra capital would not be a plus for U.S. productivity. In any case, my sense is that tracking the effect of outsourcing on productivity growth is complicated, both theoretically and empirically.

Finally, it is worth emphasizing one further point that Gordon makes about outsourcing. The illustrative example in the paper potentially tells us something about the contribution of outsourcing to productivity

growth. However, for outsourcing to contribute to the recent *acceleration* of productivity (or the earlier, mid-1990s acceleration), it would need to be the case that the pace of outsourcing has picked up. Yet, as Gordon points out, the speedup in outsourcing in recent years likely has been relatively small.

The last part of the paper turns to providing an estimate of potential output growth over the next twenty years. Gordon's estimate is about 3¼ percent a year, composed of a contribution from trend productivity growth of 2½ percentage points and a contribution from demographic and labor market factors of ¾ percentage point. The demographic and labor market analysis is quite interesting and highlights the difficult choices that must be made. I want to highlight two aspects of Gordon's estimate of potential output.

First, on the demographic side, his estimates suggest that the simple rule of thumb for translating labor productivity growth into potential output that many analysts have used for years will need to be modified. The old rule of thumb was to add 1 percentage point to productivity growth to get potential output growth. Gordon's new rule is to add ¾ percentage point.

Second, his long-run trend productivity estimate of 2½ percent a year is less than his estimate of trend productivity growth in 2001–03 of 3 percent a year. For long-run budgeting purposes, it seems wise to be conservative about the long-run prospects for productivity growth relative to the latest estimate of trend growth. In addition, Gordon's estimate of 2½ percent a year is consistent with the midrange of estimates of sustainable productivity growth that Stephen Oliner and I developed in a recent paper.[6] In that paper we analyzed the steady-state properties of a multisector growth model. This procedure provided a way to translate views about growth fundamentals (such as the rate of total factor productivity growth in the semiconductor and computer industries) into estimates of labor productivity growth. Within this framework, conservative assumptions about fundamentals generated sustainable productivity growth of about 2 percent a year, but more optimistic assumptions generated sustainable productivity growth of nearly 3 percent a year. If it were assumed that an even more substantial part of the recent productivity surprise represented

6. Oliner and Sichel (2002).

a permanent boost to growth rates, one could obtain estimates of sustainable productivity growth above 3 percent.

After all of the pieces have been put together, where does the paper's analysis leave us? Gordon's estimate of trend annual productivity growth for recent years is 3 percent, which he anticipates will drop back to 2½ percent in the longer run. The paper does not provide an explicit accounting of why the productivity trend has picked up so much, although some of the proposed explanations for the recent productivity puzzle could serve as explanations for the pickup in the trend as well. As for the recent puzzle, Gordon does a very nice job of identifying potential sources of rapid productivity growth in recent years. However, we still do not have a complete quantitative accounting, which suggests that scope remains for further work on understanding the recent behavior of productivity.

General discussion: Some on the panel discussed the need for the paper to go much further in dealing with the uncertainty in estimates of the productivity trend and in projections of future GDP trends. Steven Durlauf noted Gordon's recognition that projections were uncertain but reasoned that, to be useful for policymakers, projections must somehow quantify that uncertainty and report the error bands that surround estimates and projections. Christopher Sims found the absence of even an informal analysis of uncertainty a serious defect and related this to the fact that the estimates of trends in the variables of interest were not accompanied by any measure of uncertainty. He observed that this problem was widespread in the analysis of economic time series, where the objective is often to remove trends so as to better examine what is left. Although this approach is useful in estimating seasonality, for example, which is clearly present in many time series, it is not useful for separating business cycle fluctuations from trends since they are smoothly merged in the data. For this reason it is important to detrend time series with a methodology that treats cyclical and low-frequency variations jointly and provides some evidence on the uncertainty of both trend and cyclical components. Estimating trends with H-P filters or by connecting apparent cyclical peaks does none of this. If a methodology that provided some handle on the uncertainty were available, a useful presentation of the long-run projections in the paper would be a fan chart that illustrated the likely range of outcomes and provided some basis for discussion of those outcomes.

Aaron Tornell observed that the estimates using Kalman filters with a cyclical term were a way to isolate trend and cycle, and he questioned the merging of these with the other estimates in the paper. Because the other methods did not explicitly deal with the cyclical frequencies, the trend estimates were cyclically contaminated and the estimates of Okun's coefficient were biased. Edmund Phelps was not as bothered by the apparent cyclicality of the H-P filter estimates. He said that his own studies had convinced him that the observed swings in the unemployment rate were variations in what could loosely be called the natural rate of unemployment. On this view, substantial variability in the trend is not surprising and should not be removed from the data.

Several panelists discussed the strong productivity growth of the last three years. Jeffrey Frankel argued that part of this growth came from firms' response to the profits disappointments that followed the boom years. The urgent need to show better profits led firms to minimize employment during the initial quarters of economic recovery. But such economizing represents, at best, a one-time adjustment in employment levels. So this part of recent productivity gains is not sustainable and should not influence estimates of the productivity trend. As evidence that the productivity surge is temporary, Frankel noted that real wages, which typically rise in pace with the productivity trend, have risen only modestly over the past three years. Phelps offered two explanations for the recent productivity surge. He suggested that firms have been shedding employees who had been involved in forward-looking activities such as research and corporate planning but were not essential to current production. He added that labor hoarding by firms during the last phase of the boom could also be part of the explanation: firms had been overly optimistic when the economy initially slowed, hanging on to workers in the expectation that business would soon improve. As expectations became more realistic, they shed this hoarded labor.

References

Baily, Martin Neil. Forthcoming. *A Radical Transformation of the European Economy*. Washington: Institute for International Economics.

Basu, Susanto, John G. Fernald, and Matthew D. Shapiro. 2001. "Productivity Growth in the 1990s: Technology, Utilization, or Adjustment?" *Carnegie-Rochester Conference Series on Public Policy* 55(0): 117–65.

Basu, Susanto, and others. Forthcoming. "The Case of the Missing Productivity Growth, or, Does Information Technology Explain Why Productivity Accelerated in the United States But Not the United Kingdom?" *NBER Macroeconomics Annual 2003,* edited by Mark Gertler and Kenneth S. Rogoff. MIT Press for National Bureau of Economic Research.

Board of Trustees, Federal Old-Age and Survivors Insurance and Disability Insurance Trust Funds. 2003. *The 2003 Annual Report.* Washington: Government Printing Office.

Bresnahan, Timothy F., and Manuel Trajtenberg. 1995. "General Purpose Technologies: 'Engines of Growth'?" *Journal of Econometrics* 65(1, Special Issue): 83–108.

Bresnahan, Timothy F., Erik Brynjolfsson, and Lorin M. Hitt. 2002. "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence." *Quarterly Journal of Economics* 117(1): 339–76.

Brewster, Karin L., and Ronald R. Rindfuss. 2000. "Fertility and Women's Employment in Industrialized Countries." *Annual Review of Sociology* 26: 271–96.

Brynjolfsson, Erik, and Lorin Hitt. Forthcoming. "Computing Productivity: Firm-Level Evidence." *Review of Economics and Statistics* (November).

Brynjolfsson, Erik, and Shinkyu Yang. 1997. "The Intangible Costs and Benefits of Computer Investments: Evidence from the Financial Markets." In *Proceedings of the International Conference on Information Systems* (December). Atlanta: Association for Information Systems.

Brynjolfsson, Erik, Lorin M. Hitt, and Shinkyu Yang. 2002. "Intangible Assets: Computers and Organizational Capital." *BPEA 1:2002,* 137–81.

Clark, Peter K. 1983. "Okun's Law and Potential GNP." Working Paper. Washington: Board of Governors of the Federal Reserve System (June).

Corrado, Carol A., Charles Hulten, and Daniel E. Sichel. 2003. "Measuring Capital and Technology: An Expanded Framework." Paper presented at the CRIW/NBER conference on Measuring Capital in the New Economy, April 2001 (revised).

David, Paul A. 1990. "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox." *American Economic Review* 80(2): 355–61.

Doan, Thomas, Robert Literman, and Christopher A. Sims. 1984. "Forecasting and Conditional Projection Using Realistic Prior Distributions." *Econometric Reviews* 3: 1–100.

Eller, Jonathan W., and Robert J. Gordon. 2003. "Nesting the New-Keynesian Phillips Curve in the Mainstream Model of U.S. Inflation Dynamics." Paper presented at the Centre for Economic Policy Research Conference on "The Phillips Curve Revisited," Berlin, June 5.

Engle, Robert F., and C. W. J. Granger. 1987. "Co-Integration and Error Correction: Representation, Estimation, and Testing." *Econometrica* 55(2): 251–76.

Glassman, James. 2002. "Focus: A Feast of Possibilities." *U.S. Weekly Prospects.* J. P. Morgan Securities (November 25).

Gordon, Robert J. 1979. "The 'End-of-Expansion' Phenomenon in Short-Run Productivity Behavior." *BPEA*, *2:1979,* 447–61.

———. 1984. "Unemployment and Potential Output in the 1980s." *BPEA, 2:1984,* 537–64.

———. 1993. "The Jobless Recovery: Does It Signal a New Era of Productivity-Led Growth?" *BPEA, 1:1993,* 217–306.

———. 1997. "The Time-Varying NAIRU and Its Implications for Economic Policy." *Journal of Economic Perspectives* 11(1): 11–32.

———. 1998. "Foundations of the Goldilocks Economy: Supply Shocks and the Time-Varying NAIRU." *BPEA, 2:1998,* 297–333.

———. 2000. "Does the New Economy Measure Up to the Great Inventions of the Past?" *Journal of Economic Perspectives* 14(4): 49–74.

Hamilton, Brady E., Paul D. Sutton, and Stephanie J. Ventura. 2003. "Revised Birth and Fertility Rates for the 1990s and New Rates for Hispanic Populations, 2000 and 2001: United States." Hyattsville, Md.: National Center for Health Statistics.

Hamilton, James D. 1994. *Time Series Analysis.* Princeton University Press.

Hendry, David F., Adrian R. Pagan, and J. Denis Sargan. 1984. "Dynamic Specification." In *Handbook of Econometrics,* vol. 2, edited by Zvi Griliches and Michael D. Intriligator. Amsterdam: North Holland.

Hodrick, Robert J., and Edward C. Prescott. 1981. "Postwar U.S. Business Cycles: An Empirical Investigation." Discussion Paper 451. Minneapolis, Minn.: University of Minnesota (May).

Jorgenson, Dale W., and Kevin J. Stiroh. 2000. "Raising the Speed Limit: U.S. Economic Growth in the Information Age." *BPEA*, *1:2000,* 125–211.

Jorgenson, Dale W., Mun S. Ho, and Kevin J. Stiroh. 2002. "Projecting Productivity Growth: Lessons from the U.S. Growth Resurgence." *Federal Reserve Bank of Atlanta Economic Review* 87(3): 1–13.

Kahn, James A., and Robert W. Rich. 2003. "Tracking the New Economy: Using Growth Theory to Detect Changes in Trend Productivity." Paper presented at a conference on Technology, Productivity, and Public Policy at the Federal Reserve Bank of San Francisco, November 7–8.

Kalman, R. E. 1960. "A New Approach to Linear Filtering and Prediction Problems." *Journal of Basic Engineering, Series D* 82: 35–45.

Katz, Lawrence F., and Alan B. Krueger. 1999. "The High-Pressure U.S. Labor Market of the 1990s." *BPEA, 1:1999,* 1–65.

Kitchen, John. 2003. "A Note on the Observed Downward Bias in Real-Time Estimates of Payroll Jobs Growth in Early Expansions." Working Paper. Washington: House Budget Committee (August).

Kydland, Finn E., and Edward C. Prescott. 1990. "Business Cycles: Real Facts and a Monetary Myth." *Federal Reserve Bank of Minneapolis Quarterly Review* 14(2): 3–18.

Lawrence, Robert Z., and Charles L. Schultze. 1987. *Barriers to European Economic Growth: A Transatlantic View.* Brookings.

McKinsey Global Institute. 2001. "US Productivity Growth, 1995–2000." McKinsey & Company (October).

Nordhaus, William D. 2002a. "The Recent Recession, the Current Recovery, and Stock Prices." *BPEA 1:2002,* 199–220.

———. 2002b. "Productivity Growth and the New Economy." *BPEA, 2:2002,* 211–65.

Okun, Arthur M. 1962. "The Gap between Actual and Potential Output." *Proceedings of the American Statistical Association.* Reprinted in *Problems of the Modern Economy,* edited by Edmund S. Phelps. New York: Norton, 1965.

Oliner, Stephen D., and Daniel E. Sichel. 2000. "The Resurgence of Growth in the Late 1990s: Is Information Technology the Story?" *Journal of Economic Perspectives* 14(4): 3–22.

———. 2002. "Information Technology and Productivity: Where Are We Now and Where Are We Going?" *Federal Reserve Bank of Atlanta Economic Review* 87(3): 15–44.

O'Mahony, Mary, and Willem de Boer. 2002. "Britain's Relative Productivity Performance: Updates to 1999." London: National Institute of Economic and Social Research (March).

Phelps, Edmund S. 2003. "Economic Underperformance in Continental Europe: A Prospering Economy Runs on the Dynamism from Its Economic Institutions." Lecture presented at the Royal Institute for International Affairs, London, March 18.

Roberts, John M. 2001. "Estimates of the Productivity Trend Using Time-Varying Parameter Techniques." *Berkeley Electronic Press Journal of Macroeconomics* 1: 1–30.

Sims, Christopher A. 1974. "Output and Labor Input in Manufacturing." *BPEA, 3:1974,* 695–728.

Staiger, Douglas, James H. Stock, and Mark W. Watson. 1997. "The NAIRU, Unemployment, and Monetary Policy." *Journal of Economic Perspectives* 11(1): 33–49.

———. 2001. "Prices, Wages, and the U.S. NAIRU in the 1990s." In *The Roaring Nineties: Can Full Employment Be Sustained?* edited by Alan B. Krueger

and Robert M. Solow. New York: Russell Sage Foundation and Century Foundation Press.

Technical Panel on Assumptions and Methods. 2003. *Report to the Social Security Advisory Board.* Washington (October).

Triplett, Jack E., and Barry P. Bosworth. 2002. " 'Baumol's Disease' Has Been Cured: IT and Multifactor Productivity in U.S. Services Industries." Paper presented at a Texas A&M Conference on "The New Economy: How New? How Resilient?" April 19.

U.S. Congress, Joint Economic Committee. 2003. "A Tale of Two Employment Surveys." JEC Report. Washington (October 14).

Woodham, Douglas M. 1984. "Potential Output Growth and the Long-Term Inflation Outlook." *Federal Reserve Bank of New York Quarterly Review* 9(2): 16–23.

Yang, Shinkyu, and Erik Brynjolfsson. 2001. "Intangible Assets and Growth Accounting: Evidence from Computer Investments." Working Paper. Massachusetts Institute of Technology (May).