IAN DEW-BECKER
*Northwestern University*

ROBERT J. GORDON
*Northwestern University*

# *Where Did the Productivity Growth Go? Inflation Dynamics and the Distribution of Income*

> It's no secret that the gap between the rich and the poor has been growing, but the extent to which the richest are leaving everybody else behind is not widely known. . . . It's like chasing a speedboat with a rowboat.
>
> —Bob Herbert, *The New York Times*[1]

> There is no question that a huge gap has opened up between productivity and living standards. . . . Not since World War II have productivity and income diverged so sharply.
>
> —Louis Uchitelle, *The New York Times*[2]

THE FIRST HALF of this decade has witnessed a sharp contrast between strong output and productivity growth, on the one hand, and slow employment and median income growth, on the other. The strong growth in output combined with weak growth in hours worked has resulted in an explosion

We are above all grateful to Daniel Feenberg of the National Bureau of Economic Research for making it possible for us to access and analyze the Internal Revenue Service micro data files. We are grateful for helpful data, ideas, and insights to Michael Harper and Phyllis Otto of the Bureau of Labor Statistics, Robert Brown of the Bureau of Economic Analysis, Carol Corrado of the Board of Governors of the Federal Reserve, Mary Daly, Bradford DeLong, Julie P. Gordon, Lawrence Katz, and Mark Watson. Daniel Feenberg and Emmanuel Saez were particularly helpful in providing insights and programs that made possible our analysis of the IRS micro data files. Chris Taylor was our exemplary research assistant.

1. "The Mobility Myth," *The New York Times,* June 6, 2005, p. A23.
2. "Were the Good Old Days *That* Good?" *The New York Times,* July 3, 2005, p. C1.

in labor productivity growth, implying an underlying trend that is rising faster than in any previous subperiod of the postwar era. Yet who received the benefits? Median household income actually fell by 3.8 percent from 1999 to 2004 and grew from 1995 to 2004 at a rate of only 0.9 percent a year, a much slower rate than that of nonfarm private business (NFPB) output per hour over the same period, at 2.9 percent.[3] Similarly, the median real wage for all workers grew over 1995–2003 at an average rate of 1.4 percent a year, less than half the rate of productivity growth.[4]

The failure of the productivity growth revival to boost the real incomes and wages of the median family and the median worker calls into question the standard economic paradigm that productivity growth translates automatically into rising living standards, as in this quote from Paul Krugman:

> Productivity isn't everything, but in the long run it is almost everything. A country's ability to improve its standard of living over time depends almost entirely on its ability to raise its output per worker . . . the essential arithmetic says that long-term growth in living standards . . . depends almost entirely on productivity growth.[5]

This paper should be read in the spirit of a detective novel, whose title might be "The Case of the Missing Productivity Payoff." Our detective story is divided into two parts, macro and micro. The macro part begins with the standard identity stating that productivity growth equals real wage growth if labor's income share is constant. We examine the aggregate data that compare productivity growth with growth in alternative real wage measures, and we then ask how the post-1995 acceleration in productivity growth enters into the econometrics of price and wage dynamics. In past incarnations of dynamic Phillips curves, productivity growth has been a sideshow to the main story, if mentioned at all. The paper provides a new look at price and wage dynamics in order to assess the causes of low inflation in the decade after 1995. In light of the strong demand of the late

3. Mishel, Bernstein, and Allegretto (2005, table 1.1, p. 42). Median household income updated from 2003 to 2004 is from David Leonhardt, "U.S. Poverty Rate Was Up Last Year," *The New York Times,* August 31, 2005, p. A1.

4. Mishel, Bernstein, and Allegretto (2005, table 2.6, p. 122). Measures of real income and real wages from this source deflate nominal values by the CPI-U-RS back to 1978. (The CPI-U-RS is a supplementary, research version of the CPI-U, which revises the CPI-U back to 1978 to incorporate improvements in measurement methodology since then.)

5. Krugman (1990, p. 9).

1990s, why was inflation so low? What role did the revival of trend productivity growth play, in contrast to other beneficial supply shocks? Did the productivity growth slowdown of 1965–79 play a parallel role in creating high inflation in the 1970s? Can dynamic wage and price equations reproduce the behavior of the changes in labor's income share observed in the data?

The investigation begins with the Gordon inflation model that explains price changes by inertia, demand shocks, and supply shocks but does not include wages. A unique contribution of this paper is then to bring wages back into the study of inflation dynamics and to develop a model that includes both price and wage equations and allows each to feed back to the other. This model can capture the effect of changes in trend productivity growth on inflation, nominal wage changes, and changes in labor's income share. In dynamic simulations of the wage-price model, we find that changes in the productivity growth trend had major effects in boosting inflation during 1965–79 and in slowing it down between 1995 and 2005.

The second, micro part of the paper then examines the behavior of labor and nonlabor income as recorded in the micro data files of the Statistics of Income Division of the Internal Revenue Service (IRS) covering 1966–2001. The IRS micro files can be used to determine how much of the real income gain over various periods—for example, 1966–2001 or 1997–2001—accrued to taxpayers at the median income and to taxpayers at different percentiles of the income distribution from the 10th to the 99.99th. The IRS data have the great advantage over the more frequently used Current Population Survey (CPS) data that they allow a microscopic view of what is going on inside the upper tenth of the income distribution. We find that increasing inequality *within* the upper tenth is as important a source of growing inequality as the higher ratio of incomes in the top decile to those in the bottom decile.

Our review of the sources of increased income inequality finds that economists have placed too much emphasis on skill-biased technical change and too little on independent factors that have pushed down relative incomes at the bottom and raised them at the top. At the bottom we take a broader perspective that extends back to the 1920s, allowing us to explain the U-shaped time pattern of inequality by three U-shaped factors: the rise and fall of labor unions, the fall and rise of immigration, and the fall and rise of international trade. At the top we focus on the enormous increase in the income share of the top 1 percent and even the top 0.01 percent. Part of our analysis applies the "economics of superstars," extending

Sherwin Rosen's original argument that new technology such as CDs, the Internet, and cable TV boosts the earnings premia of the very top performers in certain fields.[6] The other part examines the large share of added income in the upper tail of the income distribution attributable to CEOs and other top corporate officers, where the facts are clear but the interpretations are controversial.

Taken together, the macro and micro parts of our detective story allow us to allocate across the income distribution the cumulative increase in real GDP attributable to the post-1995 acceleration in productivity growth. The macro analysis allows us to determine how much of the cumulative increase was broadly allocated to all income groups through lower inflation, how much went to nominal labor income, and how much went to nominal nonlabor income. The micro analysis allows us to look within the increase in real labor income to determine how much went to the top, middle, and bottom of the income distribution. In the end we find that only the top 10 percent of taxpayers had gains in real labor income per hour that kept pace with productivity growth over either the 1966–2001 or the 1997–2001 period. The micro analysis thus reconciles the paradox that median income has lagged so far behind productivity growth while labor's income share remained roughly constant, by showing that the distributional change caused median income growth to behave very differently from mean income growth.

## Issues Raised by the Macro Data on Productivity and Labor's Share of Income

We begin by examining data on the interplay between productivity growth and labor's share of income, exploiting the fact that, by definition, a constant income share of labor compensation implies that labor productivity is growing at the same percentage rate as real labor compensation per hour. How large was the post-1995 acceleration in productivity growth? Did real wages respond by accelerating in equal measure, leaving labor's income share intact, or did labor's share decline? If so, how large was the difference across alternative measures of real wage growth?

6. Rosen (1981).

We begin with a close look at the four years ending in 2005:1. Labor productivity in the NFPB sector over these four years, according to the official data as of July 2005, registered a growth rate of 3.89 percent a year. In contrast, average real hourly earnings in the total private economy increased at an annual rate of only 0.49 percent. We will show that most of this large difference can be ascribed to data and definitional issues and that the decline in labor's income share due to the remaining difference did not offset an increase in that share during 1997–2001: labor's share was about the same in 2005:1 as eight years earlier. Over a longer period going back to 1954, labor's income share has been virtually constant. These data issues include the following: data revisions, the contrast between actual and trend productivity growth, differences between the NFPB sector and the total economy, the difference between productivity growth and real compensation growth, the difference between hourly compensation and average hourly earnings, and the impact of alternative price deflators used to convert nominal wages to real wages.

### *Data Revisions*

The annual revisions of the National Income and Product Accounts (NIPA) in late July 2005 reduced the reported growth rate of real GDP over the last few years and raised the reported rate of inflation. The complementary revisions to the Bureau of Labor Statistics (BLS) productivity data reduced the annual growth rate of NFPB productivity over the recent four-year interval from 3.89 percent a year to 3.57 percent a year.

### *Actual versus Trend*

The top panel of figure 1 depicts both actual and trend productivity growth during 1950–2005. Actual productivity growth is measured as the eight-quarter rate of change in NFPB output per hour, and trend growth as a Hodrick-Prescott (HP) smoothed trend.[7] Before the recent data revisions, the HP trend reached 3.5 percent a year during 2003–04, but with the revisions it barely scrapes 3.0 percent, considerably less than actual productivity growth during 2001–05.

7. The smoothing parameter used for the HP filter is 6400 and was chosen in Gordon (2003, pp. 220–21) to avoid too much "bending" of the trend in response to the deep recession of 1981–82.

**Figure 1. Actual and Trend Changes in Output per Hour, NFPB Sector and Total Economy, Annual Growth Rates, 1950–2005**

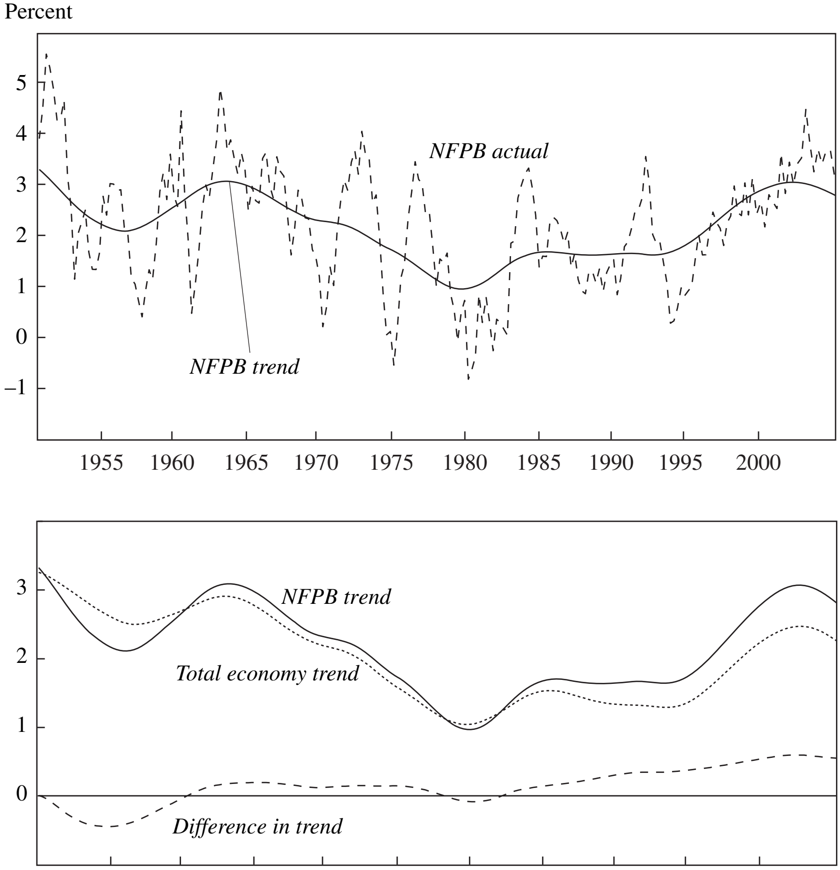


Sources: Bureau of Economic Analysis; unpublished hours data provided by Phyllis Otto, Bureau of Labor Statistics; authors' calculations.

*Total Economy versus NFPB*

The bottom panel of figure 1 again shows the HP-smoothed trend for the NFPB sector, this time juxtaposed against the similarly calculated trend for the entire economy. The latter tracks the NFPB trend quite closely until the late 1980s and then grows more slowly. It rises above

2.0 percent in late 1998, two years after the NFPB trend, and it barely reaches 2.5 percent during 2002. This panel also displays the difference between the two trends, which is very close to zero on average during 1950–85 but then begins to rise, reaching a maximum of 0.60 percent in mid-2002. Understanding this difference is important both in relation to the behavior of labor's share in the total economy and because future potential output growth depends on productivity growth in the total economy, not just that for the NFPB sector.

*Productivity, Real Compensation, and Labor's Income Share*

The lower line in the top panel of figure 2 displays the ratio of NIPA employee compensation in the total economy to NIPA net national factor income, that is, GNP minus consumption of fixed capital minus indirect business taxes. Contrary to the widespread impression that labor's share has been squeezed, there was no change in labor's share from 1997:3 to 2005:1: a substantial increase in the boom of the late 1990s was followed by a reversal in the early 2000s. Over a longer period, labor's share has fluctuated over a wider range. Two sharp increases occurred, the first in 1952–54 and a larger one in 1966–70. There were substantial fluctuations in labor's income share before 1984, but little movement has been seen since then.

The upper line in the same panel adds to NIPA employee compensation the labor component of proprietors' income, as estimated by the Economic Policy Institute.[8] Because the share of proprietors' income in total domestic income has declined over the years, and because the labor share of that income has also declined, this measure of labor's income share looks more stable. Over the entire interval, labor's share excluding proprietors' income rises from 65.1 percent in 1950:1 to 69.5 percent in 2005:1, while the share including proprietors' income barely rises at all, from 71.8 percent to 72.5 percent. Overall there seems to be little air of crisis in the data on labor's share. Especially when the labor component of proprietors' income is included, the share of labor in domestic income has floated up and down over the decades with no clear trend.

8. See Mishel, Bernstein, and Allegretto (2005, table 1.24, p. 95). The fraction of total proprietors' income that we use from this table is linearly interpolated between the selected years that they display.

**Figure 2. Labor and Nonlabor Income Shares, 1950–2005**

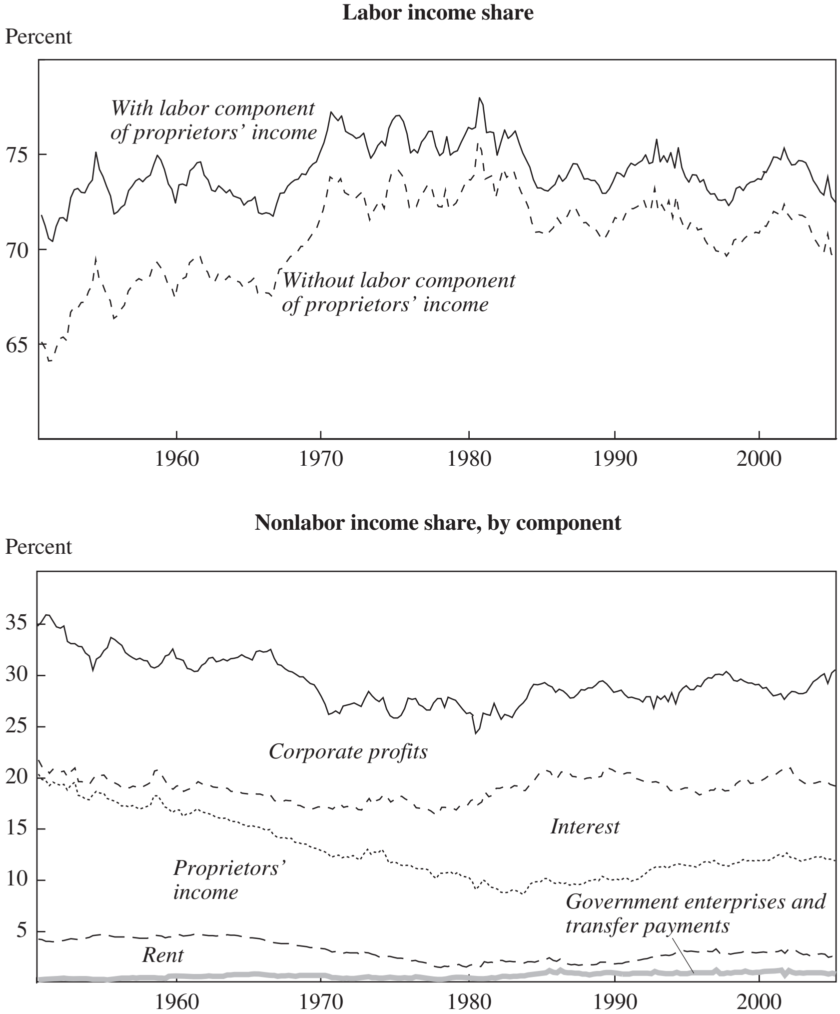


Sources: National Income and Product Accounts, Economic Policy Institute, and authors' calculations.

*Components of Nonlabor Compensation*

The bottom panel of figure 2 displays the components of the nonlabor share of domestic income, which by definition is unity minus the labor's share concept plotted as the lower line in the top panel.[9] From bottom to top (and most stable to least stable), the five components are government enterprises and transfer payments, income from rent, proprietors' income, interest, and corporate profits. The sum of rent and proprietors' income declined from 20.1 percent of income in 1950:1 to a low of 8.0 percent in 1983:3 and then gradually increased to reach 11.1 percent in 2005:1. What emerged to take its place was a huge increase in the share of interest income, which rose slowly during 1960–80 and then surged from 1980 to 1985. The share of interest income was a mere 1.4 percent in 1950:1, increased gradually to 6.5 percent in 1979:1, then sharply to 10.9 percent in 1986:2, and finally began a slow slide to 7.2 percent in 2005:1. Presumably the increase from 1950 to 1986 is due to the gradual increase in the use of debt in the economy, multiplied by the sharp increase in nominal interest rates in the late 1960s, and especially between 1978 and 1981.

Much of the current discussion of the failure of productivity gains to spill over to labor income focuses on the buoyant behavior of corporate profits in the past several years. However, the bottom panel of figure 2 shows that the share of before-tax corporate profits in nonlabor income has actually declined over time. Examining the ups and downs of the profit share over successive business cycles reveals that the cyclical low point fell from 13.1 percent in 1950:1 to 6.9 percent in 2001:3, and the cyclical high point fell from 15.4 percent in 1950:4 to 11.4 percent in 2005:1.[10] After declining over the earlier part of the postwar era, the profit share has stabilized over the past two decades and was not unusually high in 2001–05.

The top panel of table 1 reports growth rates of output per hour for the total economy, for the private business sector (farm and nonfarm, or NFPB) sector, and for the residual sector, that is, government, households,

9. In the bottom panel of figure 2 all of proprietors' income is included in nonlabor income.

10. The cyclical peak of the profit share in 1997:3 (11.6 percent) was almost identical to that in 2005:1 (11.4 percent).

**Table 1. Annual Growth Rates of Output per Hour, Compensation per Hour, and Labor Share of Income, Selected Intervals, 1954–2005[a]**

Percent

| | | | | | | | | *Addendum* | |
|---|---|---|---|---|---|---|---|---|---|
| *Item* | *1954:4–1963:2* | *1963:2–1972:1* | *1972:1–1979:1* | *1979:1–1987:4* | *1987:4–1997:4* | *1997:4–2005:1* | *Entire period, 1954:4–2005:1* | *1997:4–2001:1* | *2001:1–2005:1* |
| *Output per hour* | | | | | | | | | |
| Total economy | 2.56 | 2.12 | 1.57 | 1.36 | 1.32 | 2.48 | 1.92 | 2.10 | 2.78 |
| Private business | 2.83 | 2.54 | 1.88 | 1.60 | 1.66 | 3.01 | 2.28 | 2.54 | 3.40 |
| Farm | 4.80 | 4.62 | 4.04 | 6.81 | 3.36 | 4.03 | 4.71 | 4.96 | 3.28 |
| Nonfarm | 2.48 | 2.27 | 1.76 | 1.38 | 1.64 | 3.07 | 2.12 | 2.46 | 3.57 |
| Residual | 1.79 | 0.62 | 0.46 | 0.68 | 0.05 | 0.33 | 0.66 | 0.56 | 0.15 |
| Government | 0.66 | –0.10 | –0.23 | 0.47 | –0.21 | –0.12 | 0.08 | –0.28 | 0.02 |
| Households | 5.57 | 5.31 | 3.95 | 3.53 | 4.57 | 2.70 | 4.44 | 8.92 | –2.35 |
| Institutions | 0.29 | 1.61 | 0.80 | 0.34 | 0.01 | 0.20 | 0.57 | 0.18 | 0.22 |
| *Real compensation per hour* | | | | | | | | | |
| Total economy | 2.86 | 2.63 | 1.51 | 1.24 | 1.07 | 2.68 | 2.12 | 3.57 | 1.95 |
| Private business | 3.13 | 2.88 | 1.78 | 1.49 | 1.39 | 3.48 | 2.43 | 4.64 | 2.53 |
| Residual | 1.40 | 1.60 | 0.67 | 0.28 | –0.01 | –0.35 | 1.00 | –0.08 | –0.58 |
| *Labor share of income* | | | | | | | | | |
| Total economy | 0.31 | 0.51 | –0.06 | –0.12 | –0.25 | 0.20 | 0.20 | 1.47 | –0.83 |
| Private business | 0.30 | 0.35 | –0.09 | –0.11 | –0.27 | 0.46 | 0.16 | 2.10 | –0.87 |
| Residual | –0.39 | 0.97 | 0.21 | –0.40 | –0.06 | –0.69 | 0.33 | –0.63 | –0.73 |

Sources: Bureau of Economic Analysis, National Income and Product Accounts, table 1.3.6, and unpublished hours data provided by Phyllis Otto, Bureau of Labor Statistics.
a. All growth rates are calculated using natural logarithms.

and institutions. The first six columns show average annual growth rates between benchmark quarters,[11] the next column shows the average growth rate over the entire 1954–2005 interval, and the final two columns break the 1997–2005 period at 2001:1, in order to focus on differences between the 1997–2001 period, when the economy was enjoying its extraordinary expansion, and 2001–05, when the economy was in recession or recovering from recession. Comparison of 1954–87 with 1987–2005 shows that farm productivity growth slowed from 5.12 percent to 3.64 percent a year, and growth in the residual sector slowed from 0.90 percent to 0.17 percent a year, while NFPB productivity growth increased from 1.98 percent to 2.24 percent a year. These shifts in growth rates after 1987 account for the growing gap between the total economy and the NFPB sector in the bottom panel of figure 1.

The middle panel of table 1 reports the growth rate of real compensation per hour for the total economy, private business, and the residual sector, deflated by the NIPA deflator for each sector (*not* the GDP or the PCE deflator). Using sectoral deflators is appropriate for calculating labor's share in each sector from that sector's productivity. Over the entire period 1954–2005, differences between real wage growth and productivity growth in the total economy were minimal, with average growth of real compensation of 2.12 percent a year compared with 1.92 percent a year for productivity, implying a slight increase in labor's share of 0.20 percentage point a year over the fifty-one years. Somewhat surprisingly, in light of the frequently heard comment that labor "lost out" from the productivity growth upsurge, labor's share in the total economy actually increased at the same annual rate of 0.20 percent over 1997–2005. In the private business sector, although there have been differences in the growth rate of labor's share over shorter intervals, the broad sweep of postwar history exhibits a growth in labor's share during 1954–72, negative growth between 1972 and 1997, and then positive growth again in 1997–2005.

11. Benchmark quarters are those when the actual unemployment rate is roughly equal to the NAIRU and is declining through the NAIRU. Benchmark quarters in table 1 for the period up through 1987 are the same as or within one quarter of those used in Gordon (2003). Reflecting lower estimates of the NAIRU displayed in figure 5 below, the mid-1990s benchmark has been shifted from 1994:4 to 1997:4. The final benchmark quarter is 2005:1, the end of the data examined in table 1.

*Alternative Wage Indexes and the Role of Price Deflators*

Table 1 compared productivity growth with labor income growth measured by only a single real wage index, compensation per hour deflated by the sectoral deflators. Table 2 makes the comparison for a wider variety of real wage indexes. Compared first are three real wage indexes that use the private business deflator, of which the first is real compensation per hour in the private business economy, the same as in the middle panel of table 1. Next is the employment cost index (ECI), which is a CPI-like index of a market basket of wages that controls for shifts in mix across industries and occupations, both of which plague the hourly compensation measure.[12] Third is the BLS index of average hourly earnings (AHE) for production and nonsupervisory workers. Because the real AHE growth rate is often deflated by the CPI in official government publications, we include, in the bottom two rows of table 2, the differences in growth rates between price indexes that allow us to translate different systems of deflation for the alternative wage indexes.[13]

The growth rate of real compensation per hour is above the growth rate of private sector business productivity both over the entire 1954–2005 period and over the shorter 1997–2005 period. In contrast, for the 1954–2005 interval AHE grew more slowly than private sector business productivity by 1.13 percent. When deflated by the CPI-U, the shortfall of real AHE is −1.83 percent a year (that is, AHE using the business deflator of 1.15 percent a year minus the difference between the CPI and the business deflator of 0.70 percentage point in the bottom two lines of the table minus the productivity growth rate of 2.28 percent a year in the top line of the table).

Shifting the time interval to 1979–2005 allows the ECI to be brought into the comparison, in the far right column of the table. For this interval,

12. Because the ECI is available only back to 1978, several blank cells appear in table 2, but we are able to track the growth rate of the ECI measure of the real wage over our subintervals starting in 1979.

13. The AHE is deflated by the CPI in the *Economic Report of the President 2005,* table B-47. There are many reasons for differences in the growth rates of the PCE deflator and the CPI-U. Because of its use in indexed contracts, the CPI-U is never revised, whereas the PCE deflator is repeatedly revised to reflect improvements in methodology, which have tended to reduce the inflation rate over such periods as 1978–2000. In addition, the PCE deflator and the CPI-U incorporate different treatments of particular types of consumption, especially medical care for all years and owner-occupied housing before 1983.

**Table 2. Annual Growth Rates of Output per Hour and Real Labor Earnings, Selected Intervals, 1954–2005**

Percent

| *Item* | *1954:4–1963:2* | *1963:2–1972:1* | *1972:1–1979:1* | *1979:1–1987:4* | *1987:4–1997:4* | *1997:4–2005:1* | *Entire period, 1954:4–2005:1* | *Addendum* 1997:4–2001:1 | *Addendum* 2001:1–2005:1 | *Addendum* 1979:1–2005:1 |
|---|---|---|---|---|---|---|---|---|---|---|
| Output per hour in private business sector | 2.83 | 2.54 | 1.88 | 1.60 | 1.66 | 3.01 | 2.28 | 2.54 | 3.40 | 2.02 |
| *Earnings deflated by business deflator* | | | | | | | | | | |
| Compensation per hour | 3.13 | 2.88 | 1.78 | 1.49 | 1.39 | 3.48 | 2.43 | 4.64 | 2.53 | 2.01 |
| Employment cost index | n.a. | n.a. | n.a. | 1.60 | 1.27 | 2.20 | n.a. | 2.45 | 1.99 | 1.44 |
| Average hourly earnings | 1.85 | 2.10 | 0.09 | 0.23 | 0.72 | 1.59 | 1.15 | 2.47 | 0.87 | 0.76 |
| *Earnings deflated by PCE deflator*[a] | | | | | | | | | | |
| Compensation per hour | 3.02 | 2.85 | 1.78 | 0.73 | 0.88 | 3.08 | 2.12 | 4.07 | 2.27 | 1.44 |
| Employment cost index | n.a. | n.a. | n.a. | 0.83 | 0.75 | 1.80 | n.a. | 1.88 | 1.74 | 0.87 |
| Average hourly earnings | 1.74 | 2.07 | 0.09 | –0.53 | 0.21 | 1.19 | 0.83 | 1.90 | 0.62 | 0.19 |
| *Differences in growth rates of price indexes* | | | | | | | | | | |
| PCE deflator minus Business deflator | 0.11 | 0.03 | 0.00 | 0.76 | 0.51 | 0.40 | 0.32 | 0.57 | 0.25 | 0.57 |
| CPI-U minus PCE deflator[b] | –0.14 | 0.28 | 0.60 | 0.68 | 0.42 | 0.45 | 0.38 | 0.73 | 0.22 | 0.42 |

Source: Bureau of Economic Analysis and Bureau of Labor Statistics data.
a. PCE, personal consumption expenditure.
b. CPI-U, consumer price index for all urban consumers.
n.a., not available.

private sector compensation per hour is exactly equal to private business sector productivity growth, while the shortfall for the ECI deflated by the private business deflator is −0.58 percent, that for the ECI deflated by the PCE deflator is −1.13 percent, and that for the ECI deflated by the CPI-U is −1.55 percent.

What accounts for these discrepancies? Katherine Abraham, James Speltzer, and Jay Stewart examined the discrepancy between AHE and a series very similar to compensation per hour and found that much of the difference was explained by the fact that AHE covers only production and nonsupervisory workers. Apparently workers not covered by AHE are seeing much faster wage growth than covered workers. The difference between growth in real compensation per hour and AHE is one aspect of the distinction between median and mean income growth that is an overriding theme of this paper.

## The Effect of Changes in the Productivity Trend in a Model of Inflation Dynamics

We saw in the previous section that labor's income share has remained roughly stable over most of the postwar period, and in particular was not significantly lower in early 2005 than in the early stages of the productivity revival in mid-1997. This stability in labor's share implies that the post-1995 increase in trend productivity growth exhibited in figure 1 must have created some combination of slower inflation and faster nominal wage growth. This section provides a new analysis of the effects of productivity growth on the inflation rate, using and extending the longstanding model of inflation dynamics developed by Gordon in the late 1970s and last reported upon at a Brookings Panel conference seven years ago.[14] Did the increase in trend productivity growth in 1995–2005 reduce inflation, thus easing the job of the Federal Reserve? And did the decrease in trend productivity growth in 1965–79 raise the inflation rate and thus complicate the Federal Reserve's job in an economy already buffeted by oil and import price shocks and by the destabilizing effects of the Nixon price controls? This section begins by presenting the background and specification of the dynamic inflation model, including the role of changes in the productivity

14. Gordon (1998).

growth trend, and then discusses several closely related issues in the literature on inflation dynamics that have arisen recently, including more parsimonious specifications and shifts in parameters, the slope of the Phillips curve itself among them.[15]

### *The Gordon Model of Inflation and the Role of Demand and Supply Shocks*

The inflation equation used in this paper is similar in most details to the specification developed by Gordon twenty-five years ago.[16] It builds on earlier work that combined the Friedman-Phelps natural rate hypothesis with the role of supply shocks as a source of direct shifts in the inflation rate.[17] These supply shocks can create macroeconomic externalities in a world of nominal wage rigidity. Since the mid-1990s this research has built on the work of Douglas Staiger, James Stock, and Mark Watson by incorporating time variation in the natural rate of unemployment, resulting in a time-varying non-accelerating-inflation rate of unemployment, or TV-NAIRU.[18] The model is based on a Phillips curve that has three distinguishing characteristics: the role of inertia is broadly interpreted to go beyond any specific formulation of expectations formation to include other sources of inertia, for example in wage and price contracts; the driving force from the demand side is an unemployment or output gap; and supply-shock variables appear explicitly in the inflation equation. The specification can be written in this general form:

$$p_t = a(L)p_{t-1} + b(L)D_t + c(L)z_t + e_t, \tag{1}$$

where lowercase letters designate first differences of logarithms, capital letters designate logarithms of levels, and $L$ is a polynomial in the lag operator.

The dependent variable $p_t$ is the inflation rate. Inertia is conveyed by a series of lags on the inflation rate ($p_{t-1}$). $D_t$ is an index of excess demand (normalized so that $D_t = 0$ indicates the absence of excess demand), $z_t$ is a

15. See Galí and Gertler (1999), Galí, Gertler, and Lopez-Salido (2001), and Stock and Watson (2002, 2003) for simplified new Keynesian Phillips curves. For a discussion of changing parameters see Atkeson and Ohanian (2001) and Tetlow and Ironside (2005).

16. Gordon (1982); Gordon and King (1982).

17. Gordon (1975, 1977).

18. Staiger, Stock, and Watson (1997, 2001).

vector of supply-shock variables (normalized so that $z_t = 0$ indicates an absence of supply shocks), and $e_t$ is a serially uncorrelated error term. Distinguishing features in the implementation of this model include unusually long lags on the dependent variable, and a set of supply-shock variables that are uniformly defined so that a zero value indicates no upward or downward pressure on inflation.

If the sum of the coefficients on the lagged inflation values equals unity, then there is a "natural rate" of the demand variable ($D_t^N$) consistent with a constant rate of inflation.[19] The estimation of the TV-NAIRU combines the above inflation equation, with the unemployment gap serving as the proxy for excess demand, with a second equation that explicitly allows the NAIRU to vary with time:

$$p_t = a(L)p_{t-1} + b(L)(U_t - U_t^N) + c(L)z_t + e_t \tag{2}$$

$$U_t^N = U_{t-1}^N + \nu_t, \quad E(\nu_t) = 0, \quad \text{var}(\nu_t) = \sigma^2. \tag{3}$$

In this formulation the disturbance term $\nu_t$ in the second equation is serially uncorrelated and is uncorrelated with $e_t$. When this standard deviation $\sigma = 0$, then the natural rate is constant, and when $\sigma$ is positive, the model allows the NAIRU to vary by a limited amount each quarter. If no limit were placed on the ability of the NAIRU to vary each period, the TV-NAIRU would jump up and down and soak up all the residual variation in the inflation equation 2.

The reduced-form inflation equation 2 includes the gap between the actual unemployment rate and the NAIRU, as well as the lagged dependent (inflation) variable. In addition, five variables are included that are interpreted as supply shocks (the $z_t$ variables in equations 1 and 2), namely, the change in the relative price of nonfood non-oil imports, the change in the relative price of food and energy, changes in the relative price of medical care, the change in the trend rate of productivity growth, and dummy variables for the effect of the 1971–74 Nixon-era price controls.[20] Lag lengths

19. Although the estimated sum of the coefficients on lagged inflation is usually roughly equal to unity, that sum must be constrained to be *exactly* unity for a meaningful "natural rate" of the demand variable to be calculated.

20. The relative import price variable is defined as the rate of change of the nonfood, non-oil import deflator minus the rate of change of the dependent variable. The relative food and energy variable is defined as the difference between the rates of change of the overall PCE deflator and the "core" PCE deflator. The Nixon control variables remain the same as

**Figure 3. Trend Acceleration in Nonfarm Private Business Productivity, 1950–2005[a]**

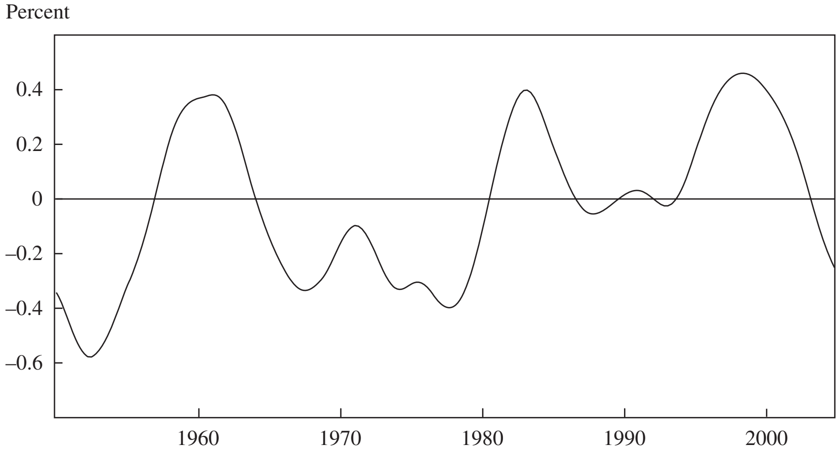


Source: Bureau of Labor Statistics data and authors' calculations.
a. Eight-quarter change in Hodrick-Prescott productivity growth trend using a smoothing parameter of 6400.

are unchanged from those originally specified by Gordon in 1982,[21] so as to allow an assessment of the robustness of this approach to twenty-five years of new data.

Besides the addition of the medical care variable, the other major change in the current inflation equation from the 1998 "Goldilocks" specification involves productivity growth, the point of departure for the current paper. Here we use the HP filter as in figure 1 to define the productivity trend and then define the acceleration or deceleration in that trend as the two-year (eight-quarter) change in the growth rate of the trend, as plotted in figure 3. Its deceleration into negative territory from 1965 to 1980 might be as important a cause of the accelerating inflation of that period as its post-1995 acceleration was as a cause of the low inflation of the late 1990s.

originally specified in Gordon (1982). Lag lengths remain as in 1982 and are shown explicitly in table 3. The medical care variable is defined as the difference between the inflation rate of the PCE deflator and the inflation rate for that deflator when medical care spending is deducted from total PCE. The productivity variable is the two-year change in the Hodrick-Prescott-filtered trend of productivity using 6400 as the smoothness parameter, as displayed in figure 4.

21. Gordon (1982).

*Coefficients in Alternative Inflation Equations*

Table 3 displays regression coefficients (or sums of coefficients), significance levels, and simulation results for our basic inflation equation and two other variants estimated over 1962:1 to 2005:2. The first data column reports results for what we call the "naïve" Phillips curve equation, which contains only the current level of the unemployment rate and four lags of the dependent variable. This equation fits the data poorly: the sum of squared residuals is 177.

The second column reports results of the 1998 "Goldilocks" version of the full specification in equation 2. Included are twenty-four lags on the dependent variable, the unemployment gap relative to the TV-NAIRU that is estimated simultaneously, and the supply-shock variables.[22] Compared with the first column, the full version in this column cuts the sum of squared residuals by almost two-thirds, from 177 to 63.

The third column reports results of our preferred specification, which incorporates both the productivity acceleration and the medical care effect and omits the deviation of productivity growth from trend. This version has better summary statistics, all of the coefficients are significant, and the simulation errors show that the equation has little drift over time and has very small mean squared error. The productivity acceleration enters with its first and fifth lags, and these coefficients sum to −1.34, indicating that an acceleration in the productivity trend reduces inflation by more than one for one. Unlike in the Goldilocks specification, which uses the deviation of actual productivity growth from trend, this variable is highly significant and shows that changes in the productivity trend have a major impact on inflation.[23]

Figure 3 gives an idea of the scale of this impact. The acceleration in the productivity trend hit its peak of 0.46 in 1999, and the effect of the variable near the peak of the last business cycle, between 1998 and 2000, would have been to lower inflation by about a half percentage point.

22. In addition to omitting the medical care effect, the Goldilocks version in this column uses the original Gordon (1982) treatment of productivity, that is, the deviation of actual productivity growth from its trend growth rate.

23. The sum of the coefficients on the two productivity acceleration terms is highly significant, with a *t* statistic of −4.07. The two individual coefficients are not significant, indicating that they convey the same information. Nevertheless, we include both rather than one or the other for expositional convenience.

**Table 3. Regressions Explaining Quarterly Changes in the PCE Deflator, 1962:1–2005:2, and Dynamic Simulation[a]**

| *Independent variable* | *Naïve specification* | *"Goldilocks" specification with productivity deviation*[b] | *Specification with productivity acceleration and medical care effect* |
|---|---|---|---|
| Constant | 1.45 | | |
| Lagged dependent variable, 1–4 lags | 0.95 | | |
| Lagged dependent variable, 1–24 lags[c] | | 1.00** | 1.00** |
| Unemployment rate | –0.20 | | |
| Unemployment gap, current and 4 lags | | –0.66** | –0.66** |
| Relative price of imports, 1–4 lags | | 0.08** | 0.07** |
| Food and energy effect, current and 4 lags | | 1.02** | 1.10** |
| Medical care effect, current and 4 lags | | | 1.57** |
| Deviation of productivity growth from trend, current and 1 lag | | –0.04 | |
| Productivity acceleration,[d] 1st lag | | | –0.41 |
| Productivity acceleration,[d] 5th lag | | | –0.93 |
| Nixon wage and price controls dummy "on" | | –1.70** | –1.52** |
| Nixon wage and price controls dummy "off" | | 1.78** | 1.97** |
| *Summary statistics:* | | | |
| Adjusted $R_2$ | 0.791 | 0.932 | 0.934 |
| Standard error of the estimate | 1.176 | 0.659 | 0.648 |
| Sum of squared residuals | 176.9 | 63.4 | 59.2 |
| *Dynamic simulation, 1995:3–2005:2*[e] | | | |
| Mean error | –4.00 | –0.64 | –0.11 |
| Root mean squared error | 4.57 | 0.83 | 0.56 |

Source: Authors' regressions.
a. All lags are quarterly. * indicates statistical significance at the 5 percent level, and ** at the 1 percent level.
b. See Gordon (1998).
c. The lagged dependent variable is entered as the four-quarter moving average for lags 1, 5, 9, 13, 17, and 21.
d. Measured as the eight-quarter change in trend productivity growth.
e. Based on regressions in which the coefficients on the lagged dependent variable are constrained to sum to unity.

Symmetrically, the deceleration in the trend hit a local minimum of –0.40 in 1978, and this added to the acceleration of inflation in the 1970s by about half a percentage point. As table 5 will show, these static results greatly underestimate the full impact of the 1995–2005 productivity growth acceleration and the 1965–78 deceleration, because they ignore dynamic feedback through the lagged dependent variable.

*The Inflation Equation: Simulation Performance*

Although most papers presenting time-series regression results display regression coefficients, significance levels, and summary statistics, few go beyond that to display the results of dynamic simulations. Yet the performance of the inflation equation is driven in large part by the role of the lagged dependent variable, making dynamic simulations the preferable method for testing. To run such simulations, we truncate the sample period ten years before the end of the data interval in 2005:2, and we use the estimated coefficients through 1995:2 to simulate the performance of the equation for those ten years, generating the lagged dependent variables endogenously.

Since the simulation has no information on the actual value of the inflation rate and no error correction mechanism, there is nothing to keep the simulated inflation rate from drifting far away from the actual rate.[24] The bottom two rows of table 3 summarize the simulation results in two statistics, the mean error over the forty-quarter simulation period and the root mean squared error. The mean error reflects the drift of the simulated value away from the actual value, so that the naïve Phillips curve on average over 1995–2005 has a predicted value of inflation that on average is fully 4.0 percentage points higher than the actual outcome. The 1998 Goldilocks specification has a much smaller mean error of –0.64 percentage point, still a substantial overprediction of inflation.[25] In contrast, our preferred specification has a minuscule mean error of –0.11. The root mean squared error of the preferred specification of 0.56 is

24. A qualification is that the TV-NAIRU used to calculate the unemployment gap after 1995 is based on data for the full period 1962–2005. This makes little difference, since the TV-NAIRU is almost constant during the 1995–2005 interval.

25. The simulation errors shown in the second column are calculated with the 1998 specification run on today's data through 2005. The simulation mean error reported in the 1998 paper (table 3, p. 315) was –0.46, but that was for a simulation period of only twenty-two quarters whereas this paper covers a more demanding forty quarters.

**Figure 4. Four-Quarter Changes in Actual, Predicted, and Simulated Values of the PCE Deflator, 1984–2005[a]**

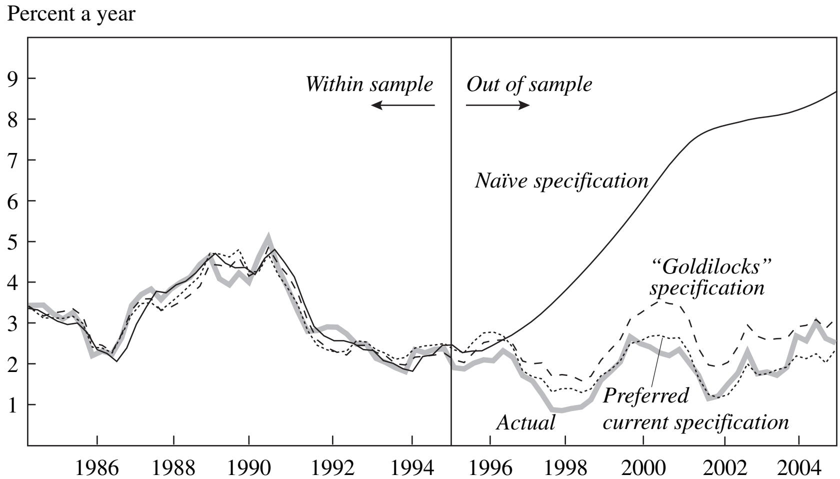


Source: Bureau of Labor Statistics data, authors' calculations.
a. Simulations are generated using the models reported in table 3.

actually substantially lower than the within-sample standard error of the estimate of 0.65.

Figure 4 displays vividly the differences among these simulation results in their ability to track the four-quarter change in the actual PCE inflation rate. Within-sample predicted values are plotted to the left of the vertical line, and the post-sample simulated values to the right. The naïve specification has no clue as to why inflation was so low in the late 1990s, and its simulated inflation rate soars to close to 9 percent by 2005. The Goldilocks specification drifts above the actual outcome, but by 2005 it is still only half a percentage point too high. The preferred specification hugs the actual values with amazing tightness.

The excellent simulation performance of the preferred specification has two important implications. First, inflation is more than simply a random walk. The supply-shock variables and the unemployment gap add a substantial amount of information beyond that from the lagged dependent variable. Second, the absence of drift in the simulations shows that the equation is stable after 1995.

Indeed, the price equation is stable not only after 1995, but across the full sample as well. A Chow test for a break at 1983:4 cannot reject the null of no break.[26] Furthermore, when interaction terms are added allowing any of the coefficients to change, none of the sums of interaction terms is significantly different from zero, except for the food and energy effect.

Andrew Atkeson and Lee Ohanian have claimed that the classic Phillips relationship between inflation and unemployment no longer holds.[27] Their conclusion, however, is entirely dependent on using a random walk to predict inflation. The significance of our coefficients, the performance of our simulations, and the stability of our model over time are at odds with their claim of instability and a structural shift in the Phillips curve. In our preferred specification the estimated change in the Phillips curve slope is not even remotely significant.

*Estimating the TV-NAIRU*

The TV-NAIRU is estimated in equation 3 simultaneously with the inflation equation 2. In the estimation process, the coefficients are forced to sum to unity. For each set of variables, there is a different TV-NAIRU. For instance, when supply-shock variables are omitted, the TV-NAIRU soars to 8 percent in the mid-1970s, since this is the only way the inflation equation can "explain" why inflation was so high in that decade. However, when the inflation equation includes the full set of supply shocks, the TV-NAIRU is quite stable.

As explained above, the NAIRU can be either so smooth as to be a constant, or so jumpy as to explain all the residual variation in the inflation equation. Rather than estimate the gain ratio for the Kalman smoother, either through a maximum likelihood estimate or by using the Stock-Watson median unbiased estimator, we impose a gain ratio of 0.0125.[28] This value was chosen as a compromise that would allow the NAIRU to vary over time yet also remove all negative serial correlation.[29] The TV-NAIRU series associated with our basic inflation equation for the PCE

26. The value of the $F(27,114)$ ratio based on estimating separate equations for 1962:1–1983:4 and 1984:1–2005:2 is 1.25, compared with a 5 percent critical F value of 1.55.

27. Atkeson and Ohanian (2001). See also Tetlow and Ironside (2005) for an analysis of changes in the slope used in the Federal Reserve's FRB/US model.

28. See Stock and Watson (1996). Specifically, we used a Quandt likelihood ratio statistic and drew our estimate of the gain ratio from their table 3.

29. We reject negative serial correlation in the TV-NAIRU, because the basic idea of the NAIRU is to reflect the gradual evolution of frictions in labor and product markets. For a further discussion of the smoothness issue, see Gordon (1998, pp. 311–12).

**Figure 5. Time-Varying NAIRU, "Goldilocks" and Current Specifications, 1962–2005**

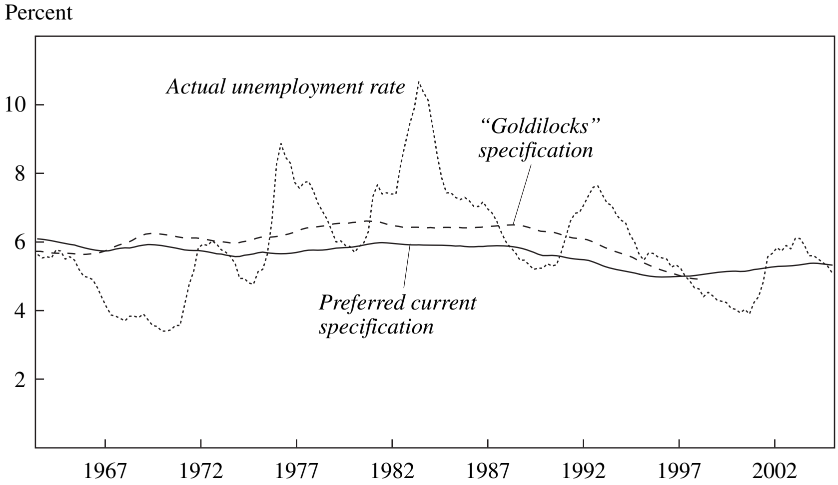


Source: Bureau of Labor Statistics data, authors' calculations.

deflator is displayed in figure 5. It remains within a narrow band between 1962 and 1988 but then drifts downward until it reaches 5 percent in 1995. Thus we concur with the general consensus that the TV-NAIRU is currently roughly in the vicinity of 5.0 percent.[30] For historical continuity, figure 5 also displays the TV-NAIRU that was estimated for the PCE deflator in Gordon's 1998 paper.[31] Our current specification yields a TV-NAIRU that is about half a percentage point below the 1998 Goldilocks specification for most of the sample period, but the 1998 version of the TV-NAIRU declines more rapidly in the mid-1990s and is virtually the same as our current version in 1997–98.

## Adding the Wage Equation and Closing the Model

It has long been recognized that any factors that affect prices may also affect wages. This can be supported from a wage aspiration framework, from the basic supply-shock perspective set out above, or from a

30. The standard error for the estimate of the NAIRU is 0.92.
31. Gordon (1998, figure 1, p. 312).

purely statistical argument.[32] Our baseline wage equation is therefore estimated with identical explanatory variables as the equation for prices. We then add a feedback term that allows wages and prices to interact dynamically.

The dependent variable in the wage equations, rather than simply compensation per hour, is the change in trend unit labor cost (TULC):

$$(4)\qquad (w - \theta^*)_t = d(L)(w - \theta^*)_{t-1} + f(L)(U_t - U_t^N) + g(L)z_t + e_t.$$

This is identical to the inflation equation 2 above except that we have replaced the inflation rate ($p$) with the change in TULC, that is, the change in nominal compensation per hour ($w$) minus the change in trend productivity growth ($\theta^*$). One reason for using the change in TULC rather than actual unit labor cost is that it removes the added variance associated with highly volatile changes in measured productivity growth. Throughout this paper we focus on trend productivity growth; it is used as an explanatory variable in tables 3 and 4 and is subtracted from changes in wages to form the dependent variable (TULC) of the wage equations.

*Closing the Model*

How was productivity growth during its post-1995 revival divided between slower inflation and faster growth in nominal compensation per hour? We start with the definition of the change in the trend labor share of income ($tls$) as the change in the nominal wage rate ($w$) minus the trend in productivity growth ($\theta^*$) minus the inflation rate ($p$):

$$(5)\qquad tls_t \equiv (w - \theta^*)_t - p_t.$$

How does the change in the trend labor share enter into the price and wage equations? An important feature of the inflation equation 2 is that wages do not matter for the determination of inflation. Similarly, in the wage equation 4, prices do not matter for the determination of wage changes. Loosening these restrictions allows us to develop symmetric price and wage equations with mutual feedback between prices and wages, which is transmitted by the change in the trend labor share.

32. Sims (1987) argued that equations with wages and prices as alternative dependent variables are simply alternative "rotations" of each other.

We begin by modifying the wage equation 4 to allow changes in TULC to depend on *both* lagged TULC changes and lagged inflation:

$$(6)\qquad (w-\theta^*)_t = d(L)(w-\theta^*)_{t-1} + h(L)p_{t-1} + f(L)(U-U^N)_t + g(L)z_t + e_t$$

$$(7)\qquad (w-\theta^*)_t = [d(L)+h(L)](w-\theta^*)_{t-1} - h(L)(w-\theta^*-p)_{t-1} + f(L)(U-U^N)_t + g(L)z_t + e_t.$$

Equation 6 is completely general in allowing *any mix* of lagged TULC change and lagged inflation to drive the evolution of TULC changes. Equation 7 is a simple algebraic rearrangement of equation 6 that adds and subtracts the "$h$" coefficients multiplied by the change in TULC $(w-\theta^*)_t$. By constraining the sum of the $d$ and $h$ coefficients, the natural rate hypothesis can be retained. Notice that the transformation brings the change in the trend labor share into the equation, as the second term multiplied by $h(L)$ is the same as $tls_t$ as defined in equation 5. An identical transformation can be applied to the price equation that adds the lagged effect of TULC as a supply-shock term and, after the same transformation, introduces the change in the trend labor share into the inflation equation[33]:

$$(8)\qquad p_t = a(L)p_{t-1} + j(L)(w-\theta^*)_{t-1} + b(L)(U_t - U_t^N) + c(L)z_t + e_t$$

$$(9)\qquad p_t = [a(L)+j(L)]p_{t-1} + j(L)(w-\theta^*-p)_{t-1} + f(L)(U_t-U_t^N) + g(L)z_t + e_t.$$

Notice that our final TULC change and inflation equations (equations 7 and 9) are completely symmetric, explaining the dependent variable with a set of lagged dependent variables, the change in the trend labor share, the output gap, and supply shocks. The only difference is that the change in the trend labor share enters with opposite signs: negative in the TULC equation and positive in the inflation equation.

33. Equation 7 is identical to equation 8 in Gordon (1998, p. 306). That paper worked out the role of changes in the trend labor share in transmitting wage impulses and price impulses back and forth between the inflation and wage change equations, but it did not develop an adequate empirical implementation of the model.

*Estimated Coefficients and Simulation Performance*

Table 4 displays results from the preferred specification (equation 9) of the inflation equation applied in the first and fourth columns to two different inflation measures, the PCE deflator and the NFPB deflator, respectively. The second and fifth columns then estimate equation 7 for the change in TULC. The biggest differences between the wage equation and the inflation equation are in the summary statistics, with a much better fit for the inflation equations. Coefficients are somewhat different in the wage equation from those in the price equations. The reaction of TULC to the unemployment gap is somewhat smaller than that of prices, and the reactions to the medical care effect and the relative price of imports are negative rather than positive, albeit insignificant.

In table 4 all the inflation and TULC equations include the change in the trend labor share, as required by equations 7 and 9 above, and this extra term has been entered with the first through the eighth lag. In all of these equations the sum of coefficients on lagged *tls* is significant and has the correct sign: positive in the inflation equations and negative in the TULC equations. The simulation errors for inflation are similar to those in the model without wage feedback, but those for TULC are noticeably better than when the lagged *tls* terms are excluded.

The coefficients are subtracted in the third and sixth columns of table 4 in order to derive an equation for the change in the trend labor share. An interesting result in the second row is that, as aggregate demand improves, as represented by a decline in the unemployment gap, *tls* is predicted to be negative as the extra demand boosts prices more than wages. This is nothing more than the famous result of the countercyclical real wage (or negatively sloped labor demand curve) debated in the late 1930s by Keynes and his critics.

Another important implication of table 4 is that the sum of coefficients on the lagged *tls* terms in the seventh row of the last column subtracts to a value of –0.87, implying that, all else equal, the growth rate of *tls* will tend toward zero through an error correction mechanism, eventually finding an equilibrium. Second, the early (first lag) effect of a productivity acceleration implies an increase in labor's share, and the late (fifth lag) effect corrects this. The long-run effect of changes in productivity on the acceleration of the labor share is extinguished by the negative coefficient on lagged *tls*.

**Table 4. Regressions Explaining Quarterly Changes in Inflation and in Trend Unit Labor Cost, with Wage-Price Feedback, 1962:1–2005:2 and Dynamic Simulation[a]**

| | *Equation using PCE deflator* | | | *Equation using NFPB sector deflator* | | |
|---|---|---|---|---|---|---|
| | *Dependent variable* | | | *Dependent variable* | | |
| *Independent variable* | *Inflation* | *Change in trend unit labor cost* | *Difference in coefficients (change in labor share)* | *Inflation* | *Change in trend unit labor cost* | *Difference in coefficients (change in labor share)* |
| Lagged dependent variable, 1–24 lags[b] | 0.99** | 1.02** | | 0.99** | 1.05** | |
| Unemployment gap, current and 4 lags | –0.56** | –0.32 | 0.24 | –0.43** | –0.30 | 0.13 |
| Relative price of imports, 1–4 lags | 0.07* | –0.09 | –0.16 | 0.11* | –0.07 | –0.18 |
| Food and energy effect, current and 4 lags | 1.14** | 0.76 | –0.38 | 0.92** | 0.82** | –0.10 |
| Medical care effect, current and 4 lags | 1.95** | –2.60 | –4.55 | 1.80* | –1.63 | –3.43 |
| PCE trend labor share,[c] 1–8 lags | 0.20** | –0.73* | –0.93 | | | |
| NFPB sector trend labor share,[c] 1–8 lags | | | | 0.24* | –0.63* | –0.87 |
| Productivity acceleration,[d] 1st lag | –0.45** | –0.25 | 0.20 | –0.19 | –0.17 | 0.02 |
| Productivity acceleration,[d] 5th lag | –0.90** | –1.17 | –0.27 | –0.94 | –1.00 | –0.06 |
| Nixon wage and price controls dummy “on” | –1.50** | –0.30 | 1.20 | –1.20 | –0.17 | 1.03 |
| Nixon wage and price controls dummy “off” | 1.95** | 0.54 | –1.41 | 2.29** | 0.60 | –1.69 |
| *Summary statistics:* | | | | | | |
| Adjusted $R_2$ | 0.936 | 0.587 | | 0.850 | 0.580 | |
| Standard error of the estimate | 0.636 | 2.020 | | 1.070 | 2.030 | |
| Sum of squared residuals | 53.8 | 544.7 | | 152.4 | 549.6 | |
| *Dynamic simulation, 1995:3–2005:2*[e] | | | | | | |
| Mean error | 0.35 | 1.77 | | 0.46 | 2.10 | |
| Root mean squared error | 0.86 | 3.43 | | 1.26 | 3.67 | |

Source: Authors' regressions.

a. All lags are quarterly. * indicates statistical significance at the 5 percent level, and ** at the 1 percent level.

b. The lagged dependent variable is entered as the four-quarter moving average for lags 1, 5, 9, 13, 17, and 21.

c. Change in the log of the trend labor share (*tls* in equation 5).

d. Measured as the eight-quarter change in trend productivity growth.

e. Based on regressions in which the coefficients on the lagged dependent variable are constrained to sum to unity.

*Counterfactual Simulations*

Although the coefficients on the productivity acceleration variable in table 4 indicate that such an acceleration should cause a shift in the level of labor's share, a better way of illustrating exactly how productivity has influenced labor's share is to calculate a dynamic simulation of the price and wage equations. We will assume first that the productivity growth slowdown of the late 1960s and 1970s never occurred. Then we will assume that the post-1995 productivity acceleration never occurred. These counterfactual simulations are calculated by using the coefficients from the regressions over the full sample to simulate price and wage changes, first with all the variables taking their actual values, and then, alternatively, with the productivity acceleration terms set to zero. The simulation that "turns off" the productivity slowdown runs from 1965:1 to 1980:1, and the simulation that "turns off" the productivity revival runs from 1995:3 to 2005:2. Recall that the 1995–2005 simulation results differ from those summarized at the bottom of tables 3 and 4, because those simulations terminated the sample period at 1995:2, whereas these use coefficients based on the entire 1962–2005 sample period and thus would be expected to have lower mean errors.

Table 5 summarizes the results of the two simulations. The top panel shows five lines of results for the NFPB deflator: the actual change, the simulated change assuming the actual behavior of the productivity growth trend acceleration variable, the counterfactual simulation that suppresses the same productivity variable to zero, the simulation error (the first line minus the second), and the counterfactual effect of the change in trend productivity growth (the second line minus the third). The middle panel shows the same for the change in TULC, and the bottom panel shows the same for the change in the trend labor share.

The two left-hand columns of table 5 summarize results for the productivity slowdown simulation of 1965–80, and the two right-hand columns do the same for the productivity revival simulation after 1995. For each simulation the first of the two columns reports the mean annual percentage rate of change over the full simulation period, whereas the second identifies any drift in the simulations by displaying the four-quarter change in the final year of each simulation.

The fourth line in each panel summarizes our findings for the simulation errors in each period. As in the simulation results presented earlier, the

**Table 5. Simulations Imposing Zero Values on Trend Productivity Change, 1965:1–1980:1 and 1995:3–2005:2**
Percent

| | *Simulation for 1965:1–1980:1* | | *Simulation for 1995:3–2005:2* | |
|---|---|---|---|---|
| *Variable and concept* | *Mean annual change* | *Change over final four quarters* | *Mean annual change* | *Change over final four quarters* |
| *NFPB sector deflator* | | | | |
| Actual | 5.34 | 9.39 | 1.55 | 2.31 |
| Factual simulation | 5.52 | 9.41 | 1.57 | 2.32 |
| Counterfactual simulation | 4.24 | 6.73 | 2.76 | 4.03 |
| Factual simulation error[a] | –0.18 | –0.02 | –0.02 | –0.01 |
| Effect of productivity change[b] | 1.28 | 2.68 | –1.19 | –1.71 |
| *Trend unit labor cost* | | | | |
| Actual | 5.47 | 8.65 | 1.82 | 3.57 |
| Factual simulation | 5.49 | 9.07 | 1.61 | 2.66 |
| Counterfactual simulation | 4.03 | 6.06 | 2.99 | 4.57 |
| Factual simulation error | –0.02 | –0.42 | 0.21 | 0.91 |
| Effect of productivity change | 1.46 | 3.01 | –1.38 | –1.91 |
| *Change in trend labor share* | | | | |
| Actual | 0.13 | –0.74 | 0.27 | 1.26 |
| Factual simulation | –0.03 | –0.34 | 0.04 | 0.34 |
| Counterfactual simulation | –0.21 | –0.67 | 0.23 | 0.54 |
| Factual simulation error | 0.16 | –0.40 | 0.23 | 0.92 |
| Effect of productivity change | 0.18 | 0.33 | –0.19 | –0.20 |

Source: Authors' calculations.
a. Actual change minus change in the factual simulation.
b. Change in the factual simulation minus change in the counterfactual simulation.

simulation errors for the NFPB deflator are very small, with no drift at all in the final four quarters of the simulation period. For the change in the TULC (the middle panel), the mean error is again modest but the error in the final year is higher, indicating an overprediction of TULC changes at the end of the 1965–80 simulation and a substantial underprediction at the end of the 1995–2005 simulation. The errors for the change in the trend labor share match those for the change in TULC, since the inflation errors are so low.

The bottom line in each panel of table 5 reports the main results of this section. The mean effect over the simulation period of the 1965–80 productivity growth slowdown was to add 1.28 percentage points to the inflation rate, 1.46 percentage points to TULC growth, and 0.18 percentage point to the change in the trend labor share. Symmetrically, the mean effect over the simulation period of the 1995–2005 productivity growth revival was to subtract 1.19 percentage points from the inflation rate, 1.38 percentage points from TULC growth, and 0.19 percentage point from the change in the trend labor share. It appears that a sustained productivity growth acceleration shifts labor's share down, and a sustained productivity growth slowdown shifts labor's share up, explaining part of the sharp jump in labor's share observed in the NIPA data for 1966–71. The second and fourth columns show that these productivity effects continue to grow, so that after fifteen years the post-1965 productivity growth slowdown had caused the inflation rate to be 2.68 percentage points higher than it would have been otherwise, and after ten years the post-1995 productivity growth revival had held down the annual inflation rate by 1.7 percentage points, with even greater effects on the change in TULC.

Overall, these results go considerably beyond the Phillips curve literature of the past decade by developing wage equations in parallel with price equations and allowing mutual feedback between them, using the change in the trend labor share as the variable by which the feedback is transmitted. We expose the wage-price model to the demanding task of staying on track in dynamic simulations, and these yield strong responses of both wage and price changes to decelerations or accelerations in productivity growth. The trend decline in productivity growth between 1965 and 1978 was a much more important contributor to high inflation in the late 1960s and the entire decade of the 1970s than is generally recognized. And the trend increase in productivity growth after 1995 was the most important single element in keeping inflation low and allowing the Federal Reserve to set short-term interest rates much lower than would otherwise have been possible.

**Changes in the Distribution of Income**

By definition, macro models deal with sums and averages and have nothing to say about the reasons why the *average* growth rate of hourly compensation is different from the *median* growth rate of hourly compensation. Further, median *annual* earnings can grow at a different rate than median *hourly* earnings if annual hours behave differently for low-paid and high-paid workers, for example when a recession like that of 1981–82 causes a sharp drop in annual hours for low-paid workers.

To address differences in average and median growth rates of compensation that occur with a shift in the distribution of income, one has to switch from macro to micro data. This part of the paper reports new results on changes in the distribution of income from Internal Revenue Service (IRS) micro data files and compares these changes with those in the Current Population Survey (CPS), the source used by most of the literature in labor economics on issues related to income inequality. Although the IRS data have numerous disadvantages, discussed below, they have the unique advantage of allowing a microscopic look at incomes within the top 10 percent of the income distribution. Whereas the CPS data are "top-coded," so that an income of $1 million in a particular year may be classified only as "greater than $100,000," the IRS data provide precise income data from tax returns for all taxpayers, no matter whether their income is $100,000, $1 million, or $10 million.

When we compare our IRS data with the CPS data for the main part of the income distribution using the conventional ratio of income in the 90th percentile to that in the 10th percentile, we obtain results for the increase in income inequality since the late 1960s that are similar in both magnitude and timing. But when we go above the 90th percentile we find significant further increases of inequality that the CPS data miss. Although most of our analysis focuses on wage and salary income, in order to highlight the comparison of mean and median growth rates of labor compensation with the growth rate of productivity, our data also allow an analysis of changes in the distribution of nonlabor income (rent, interest, dividends, and business income) and of total (labor plus nonlabor) income.

*Data Issues*

For every year between 1966 and 2001 the IRS has released data on income tax returns from over 100,000 filers; the average over the sample period is 130,000. These returns oversample filers at the very top of the distribution, so that one can study the distribution at the level of the top 0.1 percent or even the top 0.01 percent of taxpayers.[34] Our examination of these tax returns over the thirty-six-year data period is thus based on roughly 5 million observations.

Because there are minimum income requirements for filing, the data are flawed in that they omit income earned by nonfilers at the bottom of the distribution. We follow the method of Thomas Piketty and Emmanuel Saez of counting the total number of tax units in the economy by adding the total number of married couples and nonmarried adults.[35] Total tax units and total returns filed are reported in table 6, where we find that tax returns have consistently accounted for over 92 percent of tax units. Given that those who do not file necessarily have very little income and account for only 5 to 10 percent of the population, the IRS micro data allow us to obtain a very complete record of incomes actually earned. Table 6 also shows how many tax units reported wage income each year and how many hours were worked per tax unit on average. We use the hours history to illustrate general trends in hours worked, and subsequently to compare growth in IRS real compensation per hour with growth in output per hour.[36]

Income is not always faithfully reported to the IRS. Every year the Bureau of Economic Analysis (BEA) publishes data comparing its estimates of income that should be reported to the IRS with what is actually reported on tax returns. The gap between the IRS and BEA measures of adjusted gross income (AGI) ranges between 8.4 and 14.4 percent of the latter. For wages, because nearly all wage earners file tax returns, and because their wages are reported by their employers, this gap is never

34. The oversampling is extreme at the very top, where in every year between 1966 and 2001 between 3,000 and 3,500 returns are sampled for the top 0.01 percent, representing about 40 percent of those returns for 1966 and about 23 percent for 2001.

35. Piketty and Saez (2003); these authors claim that the number of married couples who file individually is insignificant.

36. Juhn, Murphy, and Topel (2002) find that many former income earners have dropped out of the labor force and thus appear neither in wage data nor in our IRS tax data.

**Table 6. Tax Units, Federal Tax Returns Filed, and Hours Worked, Selected Years, 1966–2001**

| *Year* | *No. of tax units*[a] *(thousands)* | *No. of tax returns (thousands)* | *Share of tax units filing returns (percent)* | *Share of filed returns with wages (percent)* | *Hours worked (billions)* | *Average hours worked per return filed* |
|---|---|---|---|---|---|---|
| 1966 | 75,831 | 70,160 | 92.5 | 88.5 | 154 | 2,472.4 |
| 1972 | 83,670 | 77,573 | 92.7 | 90.2 | 161 | 2,301.4 |
| 1979 | 97,457 | 92,694 | 95.1 | 89.8 | 185 | 2,229.1 |
| 1987 | 112,640 | 106,996 | 95.0 | 85.0 | 207 | 2,276.1 |
| 1997 | 129,532 | 122,422 | 94.5 | 85.3 | 244 | 2,339.4 |
| 2001 | 137,088 | 130,255 | 95.0 | 85.4 | 254 | 2,287.3 |

Sources: Piketty and Saez (2003); Internal Revenue Service and Bureau of Economic Analysis data; Phyllis Otto, Bureau of Labor Statistics; authors' calculations. Reprinted with permission.

a. Estimated as sum of married men, divorced and widowed men and women, and single men and women aged twenty and over, from *Historical Abstracts of the United States* and *Statistical Abstract of the United States*.

greater than 6 percent.[37] We make no adjustments to wages for misreporting but instead simply assume that misreporting is equally distributed across income levels.[38] Despite the problem of misreporting, our IRS data have the advantage that most income is solidly linked to W-2, 1098, and 1099 forms, and these data are not subject to the recall bias that plagues such sample surveys as the CPS.

### *Changes in the Distribution of Wage Income in IRS and CPS Data*

Our research on the income distribution can be viewed as complementary to that of David Autor, Lawrence Katz, and Melissa Kearney, who provide an extensive critical survey of the large labor economics literature on the sources of increased inequality, as well as their own new results from analysis of the CPS data.[39] The top panel of figure 6 tracks the log ratio of income in the 90th to that in the 10th percentile (P90/P10) from 1966 to 2001 in three data sets: the March CPS, the alternative May "Outgoing Rotation Group" (MORG) samples (both from Autor, Katz, and Kearney), and our IRS data.[40] Both CPS measures compare wages only for people actually working, whereas our set looks at total income for the year. People who work only part of the year may report very low wages for the full year even if their wages when working were much higher, and these people will have very low wages in the IRS data but relatively higher wage income in the CPS data. The average IRS annual income at the 10th percentile was only about $4,000 in 2001; a full-time minimum wage worker would have earned roughly $10,000 in that year.

The bottom panel of figure 6 again shows the P90/P10 ratio over time, this time expressed as an index number with the natural log equal to zero in 1973. The three measures show surprisingly similar changes, although there are subtle differences. In the IRS series, all of the increase after 1980 has occurred by 1989, whereas the March CPS series increases over a longer period, rising from zero to 0.15 during 1980–89 and then to 0.27 during 1989–2001. The pattern in the MORG data is almost identical, with an increase from 0.01 in 1980 to 0.16 in 1989 to 0.27 in 2001. The fact that

37. Park (2002, table 3).

38. One is tempted to assume that misreporting is more prevalent among the rich, who have the means with which to do it legally and the incentive to do it illegally. If this is so, our estimates of top income shares can be viewed as a lower bound.

39. Autor, Katz, and Kearney (2005).

40. We appreciate the help of Lawrence Katz of Harvard University in providing these data on the 90/10 CPS ratios.

**Figure 6. Log Ratios of 90th to 10th Labor Income Percentiles in Three Data Sets, 1966–2001**

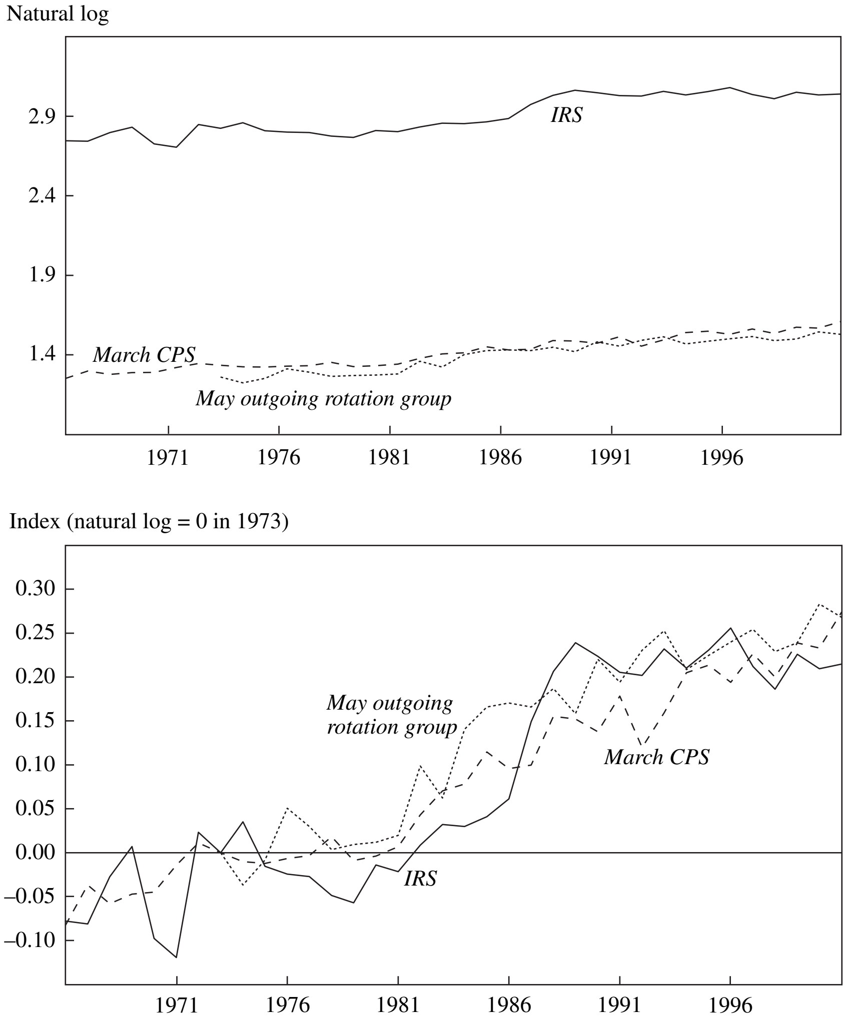


Sources: Census Bureau, March Supplement; Internal Revenue Service data; authors' calculations.

the two CPS measures continued to increase in the 1990s becomes important in distinguishing among the different explanations of increased inequality, as explored below.

Autor, Katz, and Kearney emphasize the contrast between the behavior at the bottom and at the top of the income distribution, as represented by changes in the P50/P10 and P90/P10 ratios. Their results for the CPS data are qualitatively the same as ours for the IRS data, shown in the top panel of figure 7. There was no change in the P50/P10 ratio over the full data interval. If, as this suggests, all of the increase in inequality was occurring above the 50th percentile, and not in the P50/P10 ratio, the sharp decline in the real minimum wage in the 1980s, emphasized as a major cause by David Card and John DiNardo,[41] could not have been important. In contrast to the stability of the P50/P10 ratio is the increase in the P90/P10 IRS ratio.

The bottom panel of figure 7 uses the IRS data to look within the top 10 percent of the income distribution; it shows that, whereas the log of the P90/P10 ratio increased by 0.293 over the 1966–2001 period, the log of the P99/P10 ratio increased by 0.628, and that of the P99.9/P10 ratio by 1.047. Taking antilogs to convert these into index numbers, we find that inequality by 2001 was 134, 187, and 285 percent of its 1966 level for each of the three measures, respectively. Thus the limited focus on the P90/P10 ratio in the literature based on the CPS data misses the dramatic increase in inequality within the top 10 percent.

### *The Distribution of Labor Income in the IRS Data*

We now turn to a more detailed examination of the IRS data, starting with the top panel of table 7, which summarizes the changes in the wage and salary income shares of the quantiles. Over time the shares of those in the bottom 90 percent have fallen by a total of 11 percentage points (the sum of the changes in the first four lines of the last column), while those of the top 10 percent have risen by an equivalent amount. Shown separately is the top 0.01 percent of the distribution, whose share increased by a factor of nine, from 0.2 to 1.8 percent of total wage and salary income.

Although income shares are useful for comparing relative incomes, knowing the total dollar income accruing to each quantile is more helpful for analyzing changes in welfare, and particularly for our central topic, the response of relative real incomes to the post-1995 revival in productivity

41. Card and DiNardo (2002).

**Figure 7. Selected Labor Income Percentile Ratios Using IRS Data, 1966–2001**

Index (natural log = 0 in 1966)

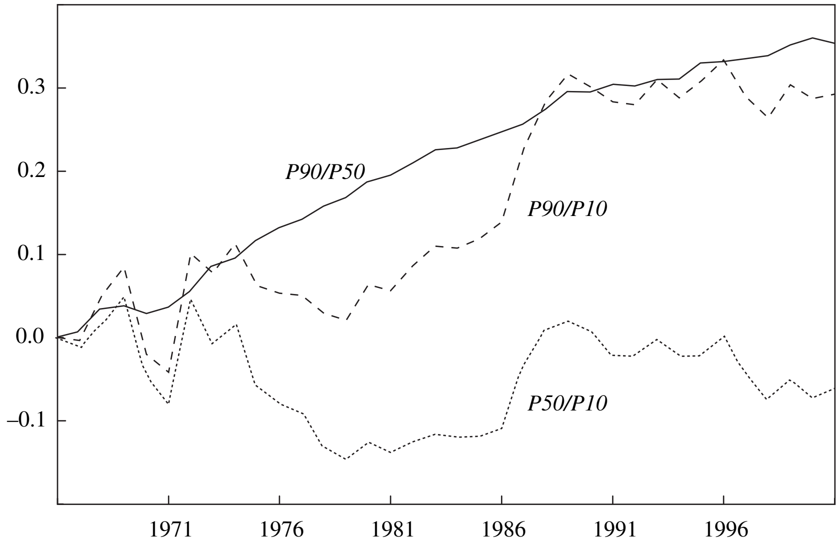


Index (natural log = 0 in 1966)

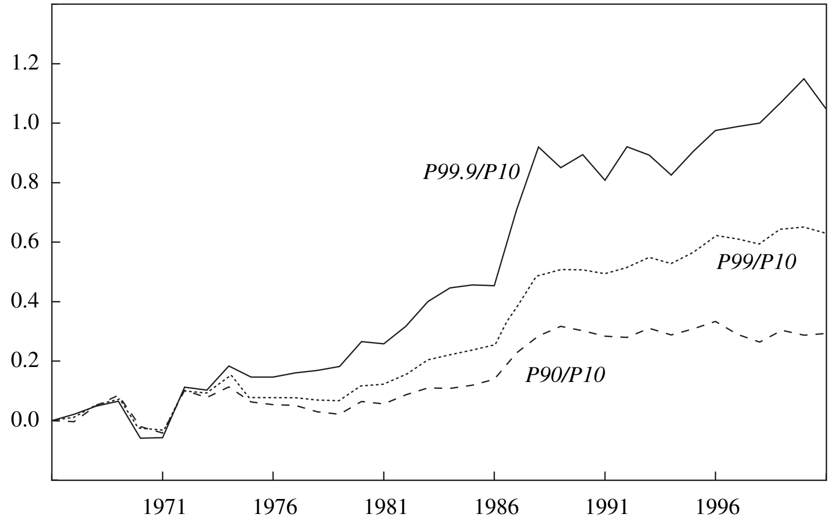


Source: Internal Revenue Service data; authors' calculations.

**Table 7. Real Wage and Salary Income and Income Shares by Quantile of the Income Distribution, Selected Years, 1966–2001**

| | | | | | | | | Share of total change (percent) | |
|---|---|---|---|---|---|---|---|---|---|
| *Percentiles* | *1966* | *1972* | *1979* | *1987* | *1997* | *2001* | *Change, 1966–2001*[a] | *1966–2001* | *1997–2001* |
| | | | | *Percent of total wage and salary income* | | | | | |
| 0–20 | 2.6 | 2.5 | 2.6 | 2.2 | 2.1 | 2.1 | –0.5 | | |
| 20–50 | 16.9 | 16.5 | 15.2 | 13.8 | 13.1 | 12.7 | –4.2 | | |
| 50–80 | 36.0 | 35.4 | 34.5 | 32.7 | 30.9 | 29.7 | –6.3 | | |
| 80–90 | 17.3 | 17.6 | 17.9 | 17.9 | 17.7 | 17.2 | –0.1 | | |
| 90–95 | 10.6 | 10.9 | 11.2 | 11.6 | 12.0 | 11.8 | 1.2 | | |
| 95–99 | 11.1 | 11.4 | 11.9 | 12.9 | 14.3 | 14.3 | 3.2 | | |
| 99–99.9 | 4.4 | 4.5 | 5.1 | 6.2 | 5.8 | 7.6 | 3.2 | | |
| 99.9–100 | 1.2 | 1.2 | 1.6 | 2.6 | 4.1 | 4.7 | 3.5 | | |
| 99.99–100 | 0.2 | 0.3 | 0.4 | 0.7 | 1.4 | 1.8 | 1.6 | | |
| | | | | *Billions of 2000 dollars* | | | | | |
| 0–20 | 43.2 | 54.2 | 68.1 | 67.5 | 78.0 | 91.9 | 48.7 | 1.7 | 1.9 |
| 20–50 | 283.0 | 360.1 | 398.3 | 422.3 | 488.6 | 569.1 | 286.1 | 10.2 | 10.8 |
| 50–80 | 603.3 | 770.3 | 900.1 | 996.2 | 1,156.4 | 1,330.2 | 726.9 | 25.9 | 23.4 |
| 80–90 | 290.5 | 382.9 | 468.2 | 546.7 | 661.5 | 771.3 | 480.8 | 17.1 | 14.8 |
| 90–95 | 178.1 | 237.2 | 293.3 | 353.7 | 448.6 | 530.0 | 351.9 | 12.5 | 11.0 |
| 95–99 | 186.0 | 248.0 | 311.2 | 394.0 | 533.5 | 640.0 | 454.0 | 16.2 | 14.3 |
| 99–99.9 | 73.3 | 98.4 | 132.3 | 188.8 | 218.6 | 339.4 | 266.1 | 9.5 | 16.2 |
| 99.9–100 | 19.5 | 27.0 | 41.1 | 80.3 | 152.5 | 210.1 | 190.6 | 6.8 | 7.7 |
| 99.99–100 | 3.8 | 5.5 | 9.4 | 22.3 | 54.2 | 82.9 | 72.1 | | |
| Total | 1,676.9 | 2,178.1 | 2,612.6 | 3,049.5 | 3,737.7 | 4,482.0 | 2,805.1 | | |

Source: Authors' calculations from Internal Revenue Service data.
a. Data in the top panel are in percentage points.

growth. The bottom panel of table 7 reports the total real wage income going to each of our selected quantiles and the change in that income between 1966 and 2001.[42] Of the total increase in real labor income of over $2.8 trillion, less than 12 percent went to the bottom half of the income distribution.[43] More of the income change accrued to the top 1 percent than to the entire lower 50 percent, and more accrued to the top 0.01 percent than to the entire bottom 20 percent. The small share going to the bottom half reflects not just growing inequality of real hourly wages, but also a smaller number of hours worked by those at the bottom.[44]

Table 7 also shows, in the far right column, how much of the increase in real compensation between 1997 and 2001 went to each quantile. Approximately the same amount went to those in the 90th through the 95th percentiles as to those in the 20th through the 50th percentiles, and about the same amount went to the bottom 80 percent as to the top 5 percent (36.1 percent and 38.2 percent, respectively). The shares of wage growth in recent years are distributed approximately the same way as those over the past thirty-six years.[45]

The top panel of table 8 shows actual wages and salaries at each threshold quantile in selected years between 1966 and 2001, as well as the skewness of the wage distribution in those years. Since skewness is unaffected by the magnitude of these values, it can be used as a consistent measure of changes in inequality over time. The most notable result here is that the median real wage has risen by only 11 percent in thirty-six years, for an average annual growth rate of 0.3 percent. Compare

42. Every inflation-adjusted number in this section is calculated using the PCE deflator, not any version of the CPI.

43. The $2.8 billion estimate for the total increase in real wage income matches nicely with the change reported by the BEA of $3.1 billion, given the decline in the percentage of BEA wages reported to the IRS.

44. Gottschalk and Danziger (2005, figures 8 and 9) show that hours worked per year for those at or below the 10th percentile are cyclically volatile compared with those at or above the 90th. Thus much of the upsurge in earnings inequality in the deep recession of 1981–82 is an artifact of annual earnings falling so much relative to hourly wages in the first group but not in the second.

45. A careful reader might note that tax units are not the same as households, and since there are an average of 1.3 tax units per household, it is possible that the average household has one tax unit that is earning very little, say, an eighteen-year-old high school senior, and one tax unit, his or her parents, that is in the upper half of the distribution and has reaped greater gains. Although this is possible, it cannot apply to most households, since the average number of tax units per household is only 1.3, and the minimum is obviously 1.

**Table 8. Real Wage and Salary Income per Tax Unit at Selected Percentiles of the Income Distribution, and Growth Rates Adjusted and Unadjusted for the Wage Compensation Share and for Hours Worked, Selected Years and Intervals, 1966–2001**

| | Percentile | | | | | | | | | Wage share of compensation (percent) | Annual growth (percent) | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| *Year* | *20th* | *50th* | *80th* | *90th* | *95th* | *99th* | *99.9th* | *99.99th* | *Skewness* | | *In hours worked* | *In productivity* |
| | *Unadjusted wage and salary income (constant 2000 dollars)* | | | | | | | | | | | |
| 1966 | 7,242 | 23,667 | 42,127 | 52,683 | 63,367 | 99,872 | 220,653 | 442,626 | 11.33 | | | |
| 1972 | 8,331 | 26,353 | 48,657 | 62,153 | 75,084 | 117,710 | 263,271 | 531,522 | 20.75 | | | |
| 1979 | 8,244 | 24,412 | 49,668 | 64,291 | 78,400 | 127,524 | 316,234 | 714,431 | 44.25 | | | |
| 1987 | 7,621 | 24,232 | 52,058 | 69,749 | 88,117 | 155,061 | 472,230 | 1,359,600 | 113.95 | | | |
| 1997 | 7,805 | 24,285 | 53,785 | 75,589 | 99,222 | 197,540 | 636,564 | 2,373,844 | 368.32 | | | |
| 2001 | 8,581 | 26,251 | 58,566 | 83,162 | 110,883 | 220,590 | 741,013 | 3,172,691 | 318.67 | | | |
| Total change, 1966–2001 (percent) | 18.5 | 10.9 | 39.0 | 57.9 | 75.0 | 120.9 | 235.8 | 616.8 | | | | |
| Average annual growth (percent) | 0.48 | 0.30 | 0.94 | 1.30 | 1.60 | 2.26 | 3.46 | 5.63 | | | | |
| | *Wage and salary income adjusted for change in share of wages in total compensation (constant 2000 dollars)*[a] | | | | | | | | | | | |
| 1966 | 7,242 | 23,667 | 42,127 | 52,683 | 63,367 | 99,872 | 220,653 | | | 90.5 | | |
| 1972 | 8,554 | 27,059 | 49,960 | 63,817 | 77,094 | 120,862 | 270,320 | | | 88.1 | | |
| 1979 | 8,916 | 26,402 | 53,717 | 69,531 | 84,790 | 137,918 | 342,009 | | | 83.7 | | |
| 1987 | 8,353 | 26,562 | 57,064 | 76,457 | 96,591 | 169,973 | 517,644 | | | 82.6 | | |
| 1997 | 8,496 | 26,436 | 58,549 | 82,285 | 108,012 | 215,039 | 692,955 | | | 83.1 | | |
| 2001 | 9,335 | 28,559 | 63,715 | 90,473 | 120,630 | 239,982 | 806,157 | | | 83.2 | | |
| Average annual growth (percent) | 0.73 | 0.54 | 1.18 | 1.55 | 1.84 | 2.50 | 3.70 | | | | | |
| | *Average annual growth in estimated wage and salary income per hour (percent)*[b] | | | | | | | | | | | |
| 1966–72 | 3.97 | 3.43 | 4.04 | 4.39 | 4.46 | 4.37 | 4.58 | | | | –1.19 | 2.17 |
| 1972–79 | 1.05 | 0.11 | 1.49 | 1.68 | 1.82 | 2.34 | 3.82 | | | | –0.46 | 1.28 |
| 1979–87 | –1.08 | –0.19 | 0.49 | 0.93 | 1.37 | 2.35 | 4.92 | | | | 0.26 | 1.43 |
| 1987–97 | –0.10 | –0.32 | –0.02 | 0.46 | 0.84 | 2.08 | 2.64 | | | | 0.27 | 1.31 |
| 1997–2001 | 2.92 | 2.49 | 2.68 | 2.93 | 3.32 | 3.31 | 4.35 | | | | –0.56 | 2.17 |
| 1966–2001 | 0.95 | 0.76 | 1.40 | 1.77 | 2.06 | 2.72 | 3.92 | | | | –0.22 | 1.57 |

Source: Authors' calculations.

a. Levels after 1996 are adjusted by the change in the aggregate wage share of compensation.

b. Change in income adjusted by the trend in hours worked.

this with average annual productivity growth of 1.57 percent over the same 1966–2001 period for the entire economy, which is somewhat slower than the 1.74 percent annual growth rate for the NFPB sector. In stark contrast, real income in the 99.9th percentile grew at over 3.4 percent a year, and that in the 99.99th percentile grew at over 5.6 percent a year. Skewness tells the same story, rising from 11 in 1966 to 319 in 2001.[46]

Since 1966, NIPA wages and salaries have made up a steadily smaller portion of NIPA total compensation as fringe benefits have risen. To correct for this, we apply the decline in the NIPA wage share of compensation equally to all percentiles. The middle panel of table 8 reports income at the threshold quantiles, adjusted for the change in the share of wages in total compensation, which itself is shown in the last column. This raises the average annual growth rate by the same amount, 0.24 percentage point, for each percentile group.

Finally, the bottom panel of table 8 shows growth rates of total compensation adjusted for growth in hours worked. In order to compare growth in total compensation with productivity growth, we need to know how many hours each tax unit worked. We make no assumptions about the distribution of the change in hours over time, but simply show what compensation growth would have been had there been no general decline in hours worked per tax unit. With these adjusted growth rates, we can compare the changes in each quantile with the 1.57 percent annual change in economy-wide productivity between 1966 and 2001. What the table shows is that *no quantile below the 90th percentile experienced growth in wages commensurate with the average rate of productivity growth.* Even the 80th percentile, after adjusting their wages upward for fringe benefits and hours, experienced real hourly compensation growth slower than average productivity growth.

Even when we look at growth in the income of individual tax units (examining a separate set of IRS panel data from 1979 to 1990), the median growth rate, after accounting for changes in hours worked per tax unit and wages as a share of compensation, rises by only 0.34 percentage point a year. This compares with a change in median income of –0.38 percent a year and economy-wide productivity growth of 1.41 percent a year

46. Measured skewness is fairly volatile, since it is heavily influenced by the top few observations, which are many orders of magnitude above the mean, but it has unambiguously risen an enormous amount over our sample.

for the 1979–90 interval. The panel data cover only a small sample of tax returns and years, but they show that, even in this alternative source of income growth that tracks individual taxpayers over time (for example, students who transition into adult jobs), the median does not keep up with productivity.

*Capital Income*

It is well known that capital (that is, nonlabor) income is less equally distributed than labor income, but did that inequality increase in the 1966–2001 period? The IRS data are ideally suited to answer this question and allow us to include five general types of nonlabor income: interest, dividends, rent, business, and pension income. We exclude capital gains income, because capital gains are excluded from NIPA personal income. Unlike wages and salaries, these income sources cannot be directly mapped to the NIPA personal income tables in any useful way, for two main reasons.[47] One is that some income covered by the IRS, such as small business income, is not included in the NIPA. The second is that there is a larger discrepancy between IRS reported income and its NIPA equivalent for nonlabor income than for wages and salaries (as discussed above).

The data on nonlabor income include many tax filers who declare losses; we exclude these returns from our data set. Further, average declared farm income is less than zero, so we completely ignore it as well. By ignoring these losses, we make the assumption that year-after-year losses are not economically meaningful but rather reflect opportunities provided by the tax system for middle- and upper-income people to shelter income from taxes. These losses do not represent what we mean by "poverty" and are economically different from the situation of those who earn only wage income and are in the bottom 20 percent of the distribution.

The top panel of table 9 reports data on total real income for selected quantiles. As one would expect, the ranking based on total income is much more concentrated than that for labor income alone, as reported in table 7. Nearly as much of the change in total real nonlabor income from 1966 to 2001 went to the top 0.1 percent as went to the bottom 50 percent.

47. The BEA does provide comparisons of BEA and IRS-equivalent measures of income, but the detailed breakdowns are not available for every year, and much of the reconciliation, especially for nonlabor income, is simply defined as "income not included in personal income," which is not helpful for the present analysis.

**Table 9. Total Income and Wage and Nonwage Income by Quantile of the Income Distribution, Selected Years, 1966–2001**

Billions of 2000 dollars except where stated otherwise[a]

| *Percentiles* | *1966* | *1972* | *1979* | *1987* | *1997* | *2001* | *Change* | *Share of total change (percent)* |
|---|---|---|---|---|---|---|---|---|
| | | | | *Total income* | | | | |
| 0–20 | 54.3 | 69.1 | 86.7 | 83.2 | 101.6 | 118.2 | 63.9 | 1.6 |
| 20–50 | 329.2 | 418.3 | 472.3 | 513.3 | 602.6 | 696.5 | 367.3 | 9.5 |
| 50–80 | 699.5 | 881.8 | 1,044.0 | 1,197.9 | 1,404.5 | 1,620.1 | 920.7 | 23.7 |
| 80–90 | 344.7 | 445.7 | 549.0 | 667.6 | 816.9 | 957.7 | 613.0 | 15.8 |
| 90–95 | 218.0 | 282.6 | 350.5 | 439.8 | 566.9 | 672.4 | 454.4 | 11.7 |
| 95–99 | 253.3 | 320.8 | 396.6 | 515.8 | 737.3 | 879.6 | 626.3 | 16.1 |
| 99–99.9 | 124.9 | 150.4 | 188.3 | 281.4 | 467.8 | 571.9 | 447.1 | 11.5 |
| 99.9–100 | 44.1 | 50.8 | 71.2 | 171.1 | 345.6 | 434.6 | 390.5 | 10.1 |
| Total | 2,068.0 | 2,619.6 | 3,158.6 | 3,870.0 | 5,043.1 | 5,951.2 | 3,883.3 | 100.0 |
| | | | | *Wage income* | | | | |
| 0–20 | 44.1 | 57.5 | 74.6 | 67.1 | 81.4 | 93.4 | 49.3 | 1.8 |
| 20–50 | 265.7 | 344.7 | 385.3 | 384.7 | 456.6 | 531.7 | 266.0 | 9.5 |
| 50–80 | 627.5 | 791.7 | 914.6 | 995.6 | 1,155.4 | 1326.3 | 698.7 | 24.9 |
| 80–90 | 309.4 | 401.5 | 488.4 | 579.2 | 692.0 | 812.3 | 503.0 | 17.9 |
| 90–95 | 187.6 | 247.2 | 304.1 | 376.9 | 474.8 | 562.1 | 374.5 | 13.4 |
| 95–99 | 177.0 | 238.8 | 298.6 | 400.6 | 543.9 | 653.9 | 477.0 | 17.0 |
| 99–99.9 | 53.4 | 77.8 | 114.6 | 176.5 | 270.5 | 327.0 | 273.6 | 9.8 |
| 99.9–100 | 12.4 | 18.2 | 32.3 | 68.9 | 126.4 | 175.2 | 162.9 | 5.8 |
| Total | 1,677.0 | 2,177.4 | 2,612.5 | 3,049.6 | 3,800.9 | 4,481.9 | 2,805.0 | 100.0 |

*(continued)*

**Table 9. Total Income and Wage and Nonwage Income by Quantile of the Income Distribution, Selected Years, 1966–2001** *(continued)*
Billions of 2000 dollars except where stated otherwise

| *Percentiles* | *1966* | *1972* | *1979* | *1987* | *1997* | *2001* | *Change* | *Share of total change (percent)* |
|---|---|---|---|---|---|---|---|---|
| | | | | *Nonwage income* | | | | |
| 0–20 | 10.2 | 11.6 | 11.9 | 15.9 | 20.2 | 24.8 | 14.6 | 1.4 |
| 20–50 | 63.5 | 74.0 | 87.1 | 128.6 | 145.9 | 164.9 | 101.3 | 9.4 |
| 50–80 | 71.9 | 90.1 | 129.6 | 202.4 | 249.5 | 294.3 | 222.3 | 20.6 |
| 80–90 | 35.3 | 44.2 | 60.6 | 88.4 | 124.9 | 145.4 | 110.1 | 10.2 |
| 90–95 | 30.0 | 35.1 | 46.3 | 62.9 | 92.3 | 110.5 | 80.5 | 7.5 |
| 95–99 | 76.3 | 82.0 | 98.0 | 115.3 | 193.3 | 225.6 | 149.2 | 13.8 |
| 99–99.9 | 71.5 | 72.9 | 73.8 | 104.9 | 197.3 | 245.0 | 173.5 | 16.1 |
| 99.9–100 | 31.8 | 32.3 | 38.7 | 102.1 | 219.2 | 259.4 | 227.6 | 21.1 |
| Total | 390.6 | 442.1 | 546.0 | 820.4 | 1,242.6 | 1,469.7 | 1,079.1 | 100.0 |
| | | | | *Nonwage income as share of total income (percent)* | | | | |
| 0–20 | 18.70 | 16.75 | 13.73 | 19.15 | 19.88 | 20.96 | 2.26 | |
| 20–50 | 19.30 | 17.69 | 18.45 | 25.05 | 24.21 | 23.67 | 4.36 | |
| 50–80 | 10.28 | 10.22 | 12.42 | 16.90 | 17.76 | 18.16 | 7.88 | |
| 80–90 | 10.24 | 9.91 | 11.03 | 13.24 | 15.29 | 15.18 | 4.94 | |
| 90–95 | 13.77 | 12.41 | 13.22 | 14.29 | 16.28 | 16.43 | 2.67 | |
| 95–99 | 30.14 | 25.57 | 24.71 | 22.36 | 26.22 | 25.65 | –4.49 | |
| 99–99.9 | 57.24 | 48.48 | 39.16 | 37.27 | 42.18 | 42.83 | –14.41 | |
| 99.9–100 | 72.00 | 63.45 | 54.33 | 59.64 | 63.43 | 59.68 | –12.32 | |

Source: Authors' calculations from Internal Revenue Service data.
a. Numbers may not sum to totals because of rounding.

The next two panels decompose income into labor and nonlabor income. Comparing shares of changes, the data for all three measures of income are roughly similar for the bottom 80 percent but then diverge sharply for the top 20 percent. The top 1 and 0.1 percent have far larger shares of the change in nonlabor income than of the change in wage income. Every other quantile takes a smaller share of the change in nonlabor income than of the change in total income.

The bottom panel of table 9 shows, not surprisingly, that as one moves up the income distribution beyond the 90th percentile, nonlabor income tends to account for a larger share of total income. Interestingly, however, nonwage income has worked its way down the income distribution over time. The top quantiles have taken most of the change in nonlabor income, but the lower quantiles, especially the 50th through the 80th and the 80th through the 90th percentiles, have seen a much larger percentage of their incomes coming from nonlabor income, whereas for the top 5 percent this proportion has declined. In 1966, 72 percent of the income of the top 0.1 percent came from nonlabor sources; by 2001 this was only 60 percent. Meanwhile, for the 50th through the 80th percentiles the share rose from 10 percent to 18 percent, nearly doubling. So there are two conflicting trends: the majority of additional nonlabor income is going to the top 10 percent of the distribution, but nonlabor income is providing a smaller share of income at the top, and a larger share in the lower quantiles.

The top and bottom panels of figure 8 show how the distribution of gains looks for the top 10 percent and the top 1 percent, respectively. The top decile tends to take about the same share of added labor and nonlabor income, but the top percentile takes a much larger share of nonwage gains. It is striking how different the bars for 1979–97 and 1997–2001 look from those for 1966–79. The share of the top 10 percent in total (labor and non-labor) real income gains ranged from 33.6 percent for 1966–79, to a much higher 59.0 percent for 1979–97, to a somewhat lower 48.6 percent for 1997–2001, averaging out at 49.4 percent for 1966–2001.

### *Lessons from the IRS Data*

Comparison of the P90/P10 and P50/P10 ratios in the IRS data confirms the basic conclusion of other authors based on CPS data that the increase in inequality since the late 1960s has been a phenomenon of the top half of the distribution, not the bottom half. But the top coding of the CPS data

**Figure 8. Shares of the Increase in Real Income Going to the Top Decile and to the Top 1 Percent, Selected Intervals, 1966–2001[a]**

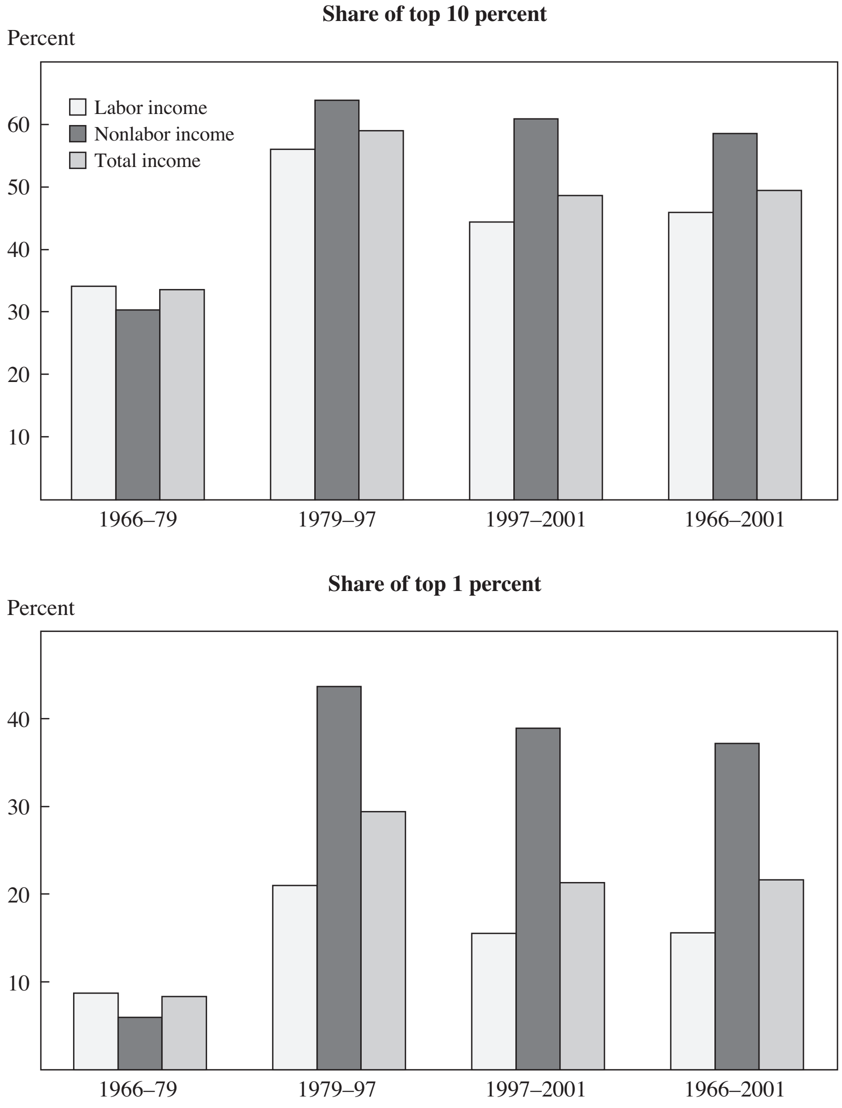


Source: Internal Revenue Service data, authors' calculations.
a. Income is measured in constant 2000 dollars.

prevents quantification of the actual dollars of real labor income earned by various strata within the top 10 percent of the distribution, or of increased skewness within the top 10 percent, which the IRS data allow.[48]

A convenient way to dramatize the role of increasing skewness at the top is to decompose the two different factors that caused fully 45.0 percent of real labor income growth from 1966 to 2001 to be earned by the top 10 percent, in sharp contrast to the 27.3 percent share in 1966. How much of this difference between 45 and 27 percent was caused by an upward movement of income above the 90th percentile relative to that in the lower percentiles, and how much was caused by increased skewness *within* the top 10 percent? To answer this question, we calculate the income of each group within the top 10 percent on the counterfactual assumption that each group's ratio to the income of the 90th percentile (P95/P90, P99/P90, and so forth) was fixed at the 1966 ratio. It is obvious from the bottom panel of figure 7 that the income level of each group above the 90th percentile would have been lower under this counterfactual, and indeed our calculation indicates that the top 10 percent would have captured 36 percent of the real income gain over 1966–2001 instead of the actual 45 percent. Thus we conclude that exactly half of the extra income gain of the top 10 percent above its original 1966 income share of 27 percent was due to increased income for the top 10 percent relative to the lower percentiles, and the rest to increased skewness *within* the top 10 percent. This second factor represents a finding that the CPS data are incapable of addressing.

Another, less widely known fact is that, as reported above, although the top of the income distribution takes most of the nonlabor income, the share of income at the 95th percentile and above that comes from nonwage sources has declined over the years, while the share for all other groups has increased.[49]

48. Gottschalk and Danziger's findings (2005, figure 16, p. 252) show enormous sensitivity to CPS measures of inequality to the method of top coding. Their P90/P10 ratio in 2002, using a base of 1979 = 100, is roughly 140 with "unadjusted top coding" and 118 with "Burkhauser top coding."

49. Our results are complementary to those of Kopczuk and Saez (2004), whose figure 9 shows that the increase in the share of total income for the top 0.01 percent over the period 1976–2000 consists almost entirely of salary and professional income rather than income from capital and capital gains. Comparing 2000 with 1929, the share of the top 0.01 percent was similar, at 3.5 percent, but a much larger share in 1929 took the form of capital income and a much smaller share the form of salaries and professional income.

Together our results thus resolve the puzzle raised at the outset: Why has growth in median real wages and real incomes lagged so far behind productivity growth when labor's share of total income has been relatively stationary? Our answer is that labor's share includes the wage and salary income of the top 10 percent, who have garnered exactly half the gains over 1966–2001. The stability of labor's share disguises a large gain in the share of that share that is going to the top 10 percent and a decline in the share going to everyone else, including the median earner.

## Causes of the Increase of Income Inequality

Our findings naturally raise questions about the interpretation of these dramatic shifts in the distribution of income that have caused median real income gains to lag so far behind productivity growth. We start with the question of whether income mobility is sufficient to mitigate the effects of rising inequality. If everyone's relative income were in constant motion, allowing each person to visit each percentile of the distribution over some span of time, there would be no cause for concern, because the cast of characters in the top 10 percent or the top 0.1 percent would be constantly changing. Next we turn to controversies in the recent literature on the causes of increased inequality and suggest a different mix of causes than has been identified in recent papers.

### *Income Mobility*

Doubts can be raised about the significance of any findings regarding income inequality that are based on a cross section of individuals who occupy different places in the income distribution from one year to the next. For one thing, income obviously depends on age, with low incomes typical during youth, higher incomes in the prime earning ages, followed by little or no labor income in retirement. An MBA student, for example, might report wage and salary income from a summer internship of $5,000 in one tax year but report income of $120,000 ten years later. Wage and salary incomes of taxpaying units fluctuate from year to year for many other reasons, including unemployment, movement in and out of the labor force in response to childbirth or illness, and fluctuations in sales commissions and bonuses in response to changes in national, local, or individual economic circumstances. Fluctuations in nonlabor income are even more likely.

How much do such factors cause our analysis above to overstate the increase in lifetime inequality? Mishel, Bernstein, and Allegretto cite a useful analogy from Joseph Schumpeter of a hotel where the quality of the rooms improves the higher the floor.[50] Is our society one in which many people over their lifetimes occupy both basement and penthouse rooms, or is it a mostly immobile society in which some remain stuck in the basement all their lives while others luxuriate permanently in the penthouse?

Evidence provided by Katherine Bradbury and Jane Katz shows clearly that there was substantial income mobility across income quintiles over decade-long periods in the 1970s, 1980s, and 1990s.[51] It would be surprising if this were not true, given the simple life-cycle factors cited above. People like the MBA student in our example account for many of those who started in the basement (the bottom 20 percent) in one year and wound up in the penthouse (the top 20 percent) just ten years later, but this basement-to-penthouse transition occurred over a full decade for only 3.3 percent of basement dwellers in 1969, 3.2 percent in 1979, and 4.3 percent in 1989. Stories like an executive's retirement, or his or her Enron- or WorldCom-like transition from the penthouse to jail, account for 5.0, 4.2, and 3.0 percent of penthouse dwellers in 1969, 1979, and 1989, respectively. Overall, the Bradbury-Katz evidence shows no increase in income mobility alongside the increase in income inequality, and indeed there were small increases in the proportion of penthouse dwellers who remained in the penthouse a decade later: from 49.1 percent in 1969, to 50.9 percent in 1979, and 53.2 percent in 1989.

In short, income mobility due to life-cycle and other reasons is a constant feature of any economy. No one is the median taxpayer or wage earner forever. The important fact about income mobility is that it takes place *independently* of the quite new phenomenon of increased skewness of the distribution of labor in the 1980s and 1990s. Not only are half of the penthouse dwellers still there a decade later, but the opulence in the penthouse keeps increasing relative to conditions in the basement.

### *Causes of Increased Income Inequality*

An enormous outpouring of literature has examined the increase in wage inequality since the 1960s. Here we cite several key contributions

50. Mishel, Bernstein, and Allegretto (2005, p. 73).
51. Bradbury and Katz (2002).

and raise some questions. We divide our discussion into three parts. The first covers the recent literature on general explanations, the most common of which is "skill-biased technical change" (SBTC), that attempt to explain the increase in such ratios as P90/P10 or P90/P50. The second concerns certain special factors operating mainly at the bottom of the income distribution. The third concerns the top of the distribution. In our view the observed changes in the income distribution reflect multiple causes, some of them independent of each other, and we reject the tendency of some analysts to argue for a particular *single* cause.

THE SBTC HYPOTHESIS. The SBTC hypothesis emerges from a simple model in which two skill classes of labor are imperfect substitutes.[52] The skilled-unskilled wage differential depends on what happens to the relative supply of the two groups of labor and on changes in the demand for skills. Often the average wages of college relative to high school graduates are used as a proxy for the skilled-unskilled differential. Because the relative quantity of college graduates has increased, particularly in the 1970s, the SBTC proponents argue that the rising skilled-unskilled wage differential must reflect a shift in employer demand toward more-skilled workers.

A prominent survey by Card and DiNardo criticizes the SBTC approach on several grounds.[53] Assuming that the dominant technical change over the past few decades has involved computers, Card and DiNardo argue that the timing is wrong. In their data all of the increase in inequality occurs during 1980–86, whereas computer technology has been developing more or less steadily over the decades, perhaps with an acceleration in the "new economy" period of the late 1990s. They also point out that the timing is wrong in relation to aggregate productivity growth, which as we have seen was slow in the 1979–97 period when inequality increased most and revived in the mid-1990s after most of the increase of inequality had already happened. Card and DiNardo much prefer an alternative explanation, namely, the decline in the real minimum wage: they find an almost perfect negative correlation between this and the increase in the P90/P10 income ratio, most of these movements being concentrated in the 1980–86 period.

Paul Beaudry and David Green also question the SBTC hypothesis, viewing it as an idea whose time has come and gone.[54] They estimate,

52. See, for instance, Katz and Murphy (1992).
53. Card and DiNardo (2002).
54. Beaudry and Green (2002).

using data for 1971–87, an SBTC equation that explains the skilled wage differential by a relative supply term and a time trend to represent technological change, and then compute the predicted value for 1988–2000. The predicted value wildly overpredicts the actual differential by about 0.35 in logs.

ALTERNATIVE HYPOTHESES. Autor, Katz, and Kearney offer a more complex, multicausal interpretation. First, they support a limited role for SBTC, taking the view that the change in the college–high school wage differential is well explained by a steady, demand-driven growth in the relative demand for college graduates overlaid with fluctuations in the relative supply of college graduates. Second, in their CPS data, as in figure 6 above, the increase in P90/P10 wage inequality is relatively steady over the 1980s and 1990s rather than concentrated only in the early 1980s when most of the drop in the real minimum wage occurred.

However, like Beaudry and Green, Autor, Katz, and Kearney also criticize a purely SBTC explanation on the ground that the increase in inequality began to plateau around 1992, whereas the new economy revival of aggregate productivity growth began around 1995. Thus they echo the skepticism of Card and DiNardo. Second, they emphasize the difference in magnitude between the P50/P10 inequality changes, which were negligible (as shown in figure 7 above), and the P90/P50 changes, which were substantial and continuous. They argue for a more articulated conception of SBTC that distinguishes among five types of job tasks, ranging from "routine manual" at the bottom to "nonroutine analytic" and "nonroutine interactive" at the top. Using occupational data, they assign different shares of these job tasks to each decile of the wage distribution and conclude that demand growth has sharply shifted toward those tasks most common in the upper three income deciles; they view this as evidence in favor of a "polarization hypothesis." However, their broad definitions of these job tasks cover substantial shares of the population and so do not explain increased skewness within the top 10 percent.

If SBTC had been a major source of the rise in inequality, we should have observed an increase in the relative wages of those most directly skilled in the development and use of computers. Yet during 1979–97 fully half of the growth in the college-noncollege wage premium can be attributed to the increased relative wage of the occupational group called "managers," and only 17 percent to the occupational groups presumably favored by SBTC (including "engineers" and "math/computer").[55] Here Europe may provide some perspective, because the increase in the ratio of CEO pay to average worker pay so evident in the United States has not occurred there. We return below to the puzzle of rising CEO pay premia.

*Inequality at the Bottom: The "Great Compression" and Its Causes*

A significant limitation of most of the SBTC literature is that it considers only the period since about 1970 and ignores the preceding fifty years. Yet the basic facts to be explained about income equality are not one but two: not only why inequality rose after the mid-1970s but also why it declined from 1929 to the mid-1970s.[56] Claudia Goldin and Robert Margo have called the flattening of the income distribution during 1930–70 the "Great Compression,"[57] and they attribute it to at least three events that fit neatly into this U-shaped pattern, all of which influence the effective labor supply curve and the bargaining power of labor: the rise and fall of unionization, the decline and recovery of immigration, and the decline and recovery in the importance of international trade and the share of imports. Union membership first rose and then declined in part because government legislation encouraged its rise in the 1930s and increasingly discouraged it in the postwar years. In addition, the invention of air conditioning facilitated the dispersion of employment into the Old Confederacy with its "right to work" laws. Unions were further weakened by the steady decline in the share of employment in manufacturing and mining, given unions' failure to organize most employees in the services sector.

Partly as a result of restrictive legislation in the 1920s, but also as a result of the Great Depression and World War II, annual immigration as a fraction of total population declined from 1.3 percent in 1914 to 0.02 percent in 1933. Immigration then remained very low until a gradual recovery

55. See Mishel, Bernstein, and Schmitt (2001, table 2.45, p. 200).

56. The most vivid representation of the U-shaped historical pattern of income inequality is Kopczuk and Saez (2004, figure 9), which shows the income share of the top 0.01 percent and its composition across labor and capital income. This share fell from 3.7 percent in 1929 to 0.6 percent in 1976 and then rose to 3.6 percent in 2000. The reason that their 3.6 percent is so much higher than our figure of 1.8 percent for 2001 in table 7 is that the Kopczuk-Saez shares for the top 0.01 percent include capital income and capital gains income, whereas table 7 reports only wage and salary income.

57. Goldin and Margo (1992).

began in the late 1960s, reaching 0.48 percent (legal and illegal immigration combined) in 2002.[58] Competition for unskilled labor arrives not only in the form of immigrants but also in the form of imports, and the decline of the import share from the 1920s to the 1950s and its subsequent recovery are basic facts of the national accounts.

Although Card and DiNardo and Autor, Katz, and Kearney raise important questions about the SBTC hypothesis, we note two others. First, inequality decreased as much from the 1920s to the 1970s as it has increased from the 1970s to now. Are we to believe that technical change over 1920–70 was "unskilled biased"? It is possible that the heyday of unionized, assembly-line manufacturing provided an abundance of repetitive jobs for high-school dropouts, but the fact that these jobs paid relatively well depended perhaps more on the strength of unions and the relative absence of immigration and imports. Second, the SBTC hypothesis fails to explain the absence of an increase of income inequality in Europe despite the free flow of technology across borders.[59]

### *Skewness at the Top: The Superstar Phenomenon*

Our analysis of the IRS data suggests that most of the shift in the income distribution has been from the bottom 90 percent to the top 5 percent, and especially to the top 1 percent. This is much too narrow a group to be consistent with a widespread benefit from SBTC. We argue in this section that two possibly independent phenomena are taking place at the top of the income distribution. The first is the increasing income premia being paid to "superstars," the subject of a brilliant analysis by the late Sherwin Rosen almost a quarter-century ago.[60] Rosen explains why a limited number of top performers in particular fields earn most of the income, and we extend his ideas to explain why the superstar premium has been increasing. We also take an explicit look at the incomes of two classes of superstars, "power celebrities" and major league athletes. A second group who earn a larger share of income at the top are CEOs and other

58. Annual immigration as a share of the U.S. population is shown in Gordon (2003, figure 5, p. 268).

59. For the latest data on the change in inequality in the United States versus the European countries, see Mishel, Bernstein, and Allegretto (2005, chapter 7). For an attempt to develop theories of how European institutions distort the path of technical change, see Acemoglu (2002).

60. Rosen (1981).

top corporate officers. Recent economic research provides the beginning of an analysis of CEO premia while leaving some important questions unanswered.

Rosen explains the extreme skewness in occupational categories dominated by superstars by particular characteristics of demand and supply. On the demand side, both market size and price per unit (the "ticket price") multiply together to form a revenue function whose convexity implies that "small differences in talent become magnified in large earnings differences, with great magnification if the earnings-talent gradient increases sharply near the top of the scale."[61] Competition remains present but does not work to drive down these differentials: "hearing a succession of mediocre singers does not add up to a single outstanding performance." On the supply side, the performer exerts the same effort whether 10 or 10,000 witness the performance.

Superstars benefit from skewness, but why has the degree of skewness increased? As Rosen recognizes, a succession of innovations going back to the phonograph, radio, and motion pictures has increased the size of audiences who can hear a given performance, thus increasing the incomes of superstars by many multiples. Thus superstars represent a particular type of SBTC that is concentrated at the very top of the income distribution, where the technological change in question is the development of compact discs, cable television, and other forms of "audience magnification." As Rosen shows, superstars represent an equilibrium phenomenon: there is no suggestion that markets do not work, and technological change feeds directly into increased premia.

A typical reaction to our use of the superstar model to explain increased skewness at the top is, "But there aren't enough superstars." This raises the question, How many superstars are there among the 13,000 IRS tax returns in the 99.99th percentile, which in 2001 accounted for $83 billion of income, with an entry threshold of $3.2 million? Here we report on the incomes of a small set of entertainment superstars and a larger group of professional athletes. But this is only the tip of the iceberg of the superstar phenomenon. Rosen himself cites examples of other entertainment superstars (comedians, classical musicians) as well as economics textbook authors, some other authors, and lawyers, and he quotes approvingly from Marshall's *Principles,* which mentions high-earning barristers, jockeys, painters, and musicians.

61. Rosen (1981, p. 846).

An annual *Forbes* feature on the "The Celebrity 100" reports estimated 2004 incomes for 100 top celebrities, including superstars in the worlds of movies, music, TV, multimedia, sports, and fiction writing.[62] The reported incomes range from $290 million for filmmaker George Lucas to a mere $1.5 million for child star Amanda Bynes. The total income accounted for is $3.1 billion, for an average annual income of $31 million, with a median income of $25 million. All but 3 of the top 100 have incomes that qualify them for the IRS 99.99th percentile. Yet this is an underestimate of the top 100 superstar incomes, because the celebrity list is chosen based not just on income but also on indicators of fame, including numbers of magazine covers, media citations, and Internet hits.

Although the top-100 celebrity list leaves unknown the total incomes of other superstars in the entertainment world, we can perform a complete census of major league baseball, football, and basketball players using data maintained by *USA Today.*[63] The 2005 total payroll for these 2,820 athletes was $7.0 billion, for an average of $2.48 million per player, a bit short of the entry level to the IRS 99.99th percentile. Here our main interest is in the contribution of these athletic salaries to the overall increase in skewness, and unfortunately only the baseball data source provides information going back more than five years. For the twenty-six major league baseball teams that have existed since 1988, total payroll has increased from $295 million to $2 billion, and the average salary has increased from $354,000 to $2,075,287. The inflation-adjusted increase was 8.9 percent a year, compared with a 6.0 percent annual rate of real increase for the IRS top 99.99th percentile between 1987 and 2001.

Together the incomes of these 100 celebrity superstars and roughly 3,000 athletes account for $10 billion in 2004–05, still well short of the $83 billion in the IRS 99.99th percentile in 2001. But this excludes other sports, such as golf and tennis, which Rosen specifically mentions as beneficiaries of media expansion, and it excludes high-earning entertainment figures below the level of the 100 celebrities. Finally, this tally excludes such celebrity lawyers as the late Johnny Cochran and David Boies and top-earning figures in such professions as management consulting and investment banking, for example Jack Meyer, who earned $25 million annually to manage the Harvard endowment.

62. *Forbes,* "The Celebrity 100," July 4, 2005, p. 102–11.

63. See asp.usatoday.com/sports/baseball/salaries/default.aspx for baseball and the corresponding sites for football and basketball.

*CEOs and Other Top Executives*

Clearly a large share of the income at the very top of the income distribution is earned by CEOs and other top corporate officers. But we treat the increasing pay premia of CEOs as different from the superstar phenomenon, in light of the puzzles that arise in its economic analysis. Here we cite other sources that document the increase in CEO pay relative to average pay, quantify total CEO compensation in a large sample of U.S. firms, and discuss alternative explanations of the increased CEO premia that doubtless contribute to the increased skewness at the top of the U.S. income distribution.

The ratio of average CEO pay to average worker pay increased from 27 in 1973 to 300 in 2000, then fell to 237 in 2001 as a result of the stock market crash.[64] Including both cash and equity compensation, the 1989–2000 increase in CEO compensation was 342 percent, which compares with 5.8 percent for the median hourly wage. A basic difficulty for any equilibrium theory to explain this jump in CEO premia is that it is primarily a phenomenon of the United States and has not happened elsewhere. The ratio of average CEO pay to the average compensation of manufacturing production workers in 2003 was 44.0 in the United States, more than double the ratio of 19.9 in thirteen other rich countries.[65]

Lucian Bebchuk and Yaniv Grinstein provide valuable data and analysis of the CEO pay phenomenon.[66] They report average pay for the top five executives in 1,500 firms: those in the S&P 500, the Mid Cap 400, and the Small Cap 600 indexes. Average pay in 2001 was $14.3 million for the CEOs and $31.9 million for the top five executives, or $6.4 million each. This is 7,500 people making $6.4 million each, more than half of the 13,000 people in the IRS 99.99th percentile, who, coincidentally, made an average of $6.4 million each in 2001. Our first inference is that most of these executives are in the IRS 99.99th percentile and that their total income of $48 billion accounts for more than half of income in that quantile in 2001.

64. Mishel, Bernstein, and Allegretto (2005, figure 2Y, p. 214). Other facts in this and the next paragraph come from the same source, pp. 212–16.

65. The discrepancy between the ratio of 44 for the cross-country comparisons and the ratio of 237 just above for the United States is explained in Mishel, Bernstein, and Allegretto (2005, p. 216) as due to inconsistent data sources.

66. Bebchuk and Grinstein (2005).

Does rising CEO pay reflect a reward for increasing firm size, rate of return, or growth in rate of return? Bebchuk and Grinstein analyze data from 14,154 firms and show that the compensation of the top five executives increased during 1993–2003 by 76 percent more than their regressions of compensation on these explanatory factors can explain. They calculate that the ratio of top-five compensation to total profits for a large sample of about 20,000 firms rose from 5.0 percent in 1993–95 to 12.8 percent in 2000–02.

What Bebchuk and Grinstein call the "arm's-length bargaining perspective" explains increased CEO pay by demand and supply factors, where demand depends on the value to corporations of executive services, and supply depends on alternatives for CEOs in other occupations and the nonpecuniary aspects of the job (stress, legal risk, and so forth). Bebchuk and Grinstein emerge unconvinced by this model, arguing that the stock market boom of the 1990s should have increased CEO premia only temporarily, not permanently, and that incomes in alternative occupations have increased much more slowly than for CEOs. They also reject the view that corporate compensation committees were ignorant of the fact that stock options imposed costs on shareholders.

Their alternative, "managerial power" perspective veers further from equilibrium economics. In this view directors do not seek to get the best deal for shareholders. The only constraint on CEOs paying themselves unlimited amounts is the "outrage constraint," which weakened in the 1990s as rising stock markets pacified shareholders. This approach comes close to being observationally equivalent to saying that CEO pay depends on stock market valuation, although Bebchuk and Grinstein's own regressions suggest that CEO pay has far outpaced such valuations. These regressions would be consistent with an influence of stock market prices if CEO pay is *correlated* with stock prices while still outpacing them.

Although their paper is inconclusive about the merits of an "arm's-length bargaining model" versus a "managerial power model," we propose a variant of the latter called the "scratch-my-back model," which posits an exclusive class of CEOs who determine each other's pay subject to relatively few market constraints. This image is conveyed by the following newspaper account:

> The compensation committee talks to an outside consultant who has surveys that you could drive a truck through and pay anything you want to pay. . . . The outside consultant talks to the HR [human resources] vice president, who talks

> to the CEO. The CEO says what he'd like to receive. It gets to the HR person who tells the outside consultant. And it pretty well works out that the CEO gets what he's implied he thinks he deserves, so he will be respected by his peers.[67]

*The Economist* christens this phenomenon the "Lake Wobegon effect," after Garrison Keillor's mythical Midwestern town where "all the children are above average." No corporate compensation committee wants to pay the average; rather they all want to pay above average, and "so bosses' pay spirals upwards."[68]

The basic data reviewed here and the controls applied in the Bebchuk-Grinstein regressions suggest that top executive compensation has spiraled up at about the same rate as the compensation of baseball players. Together the well-understood phenomenon of superstars and the puzzling case of CEOs clearly explain most of the large increase of compensation in the 99.99th percentile relative to the 90th. A broader interpretation of technology that includes new media inventions seems crucial in explaining the rising skewness of superstar income, whereas some role for stock market valuations may help to explain the CEO puzzle, including the difference between the United States and Europe in CEO pay.

## Conclusion and Further Research Agenda

This paper started as a detective story in search of the missing productivity payoff. The macro part of our investigation conducted a detailed search to locate the effects of the post-1995 productivity revival, with a parallel search for the effects of the post-1965 productivity growth slowdown. The micro part of our research used IRS data to shed light inside the top 99.9th and 99.99th percentiles and linked increased skewness at the top both to the economics of superstars and to the less well understood phenomenon of escalating CEO pay premia.

### *The Surprising Micro Conclusion*

A basic tenet of economic science is that productivity is the seed that creates the flower of a nation's standard of living. But our results raise

67. Gretchen Morgenstern, "How to Slow Runaway Executive Pay," *The New York Times* (business section), October 23, 2005, p. 1.

68. *The Economist,* "Where's the Stick?" October 11, 2003, p. 13.

doubts. Our most surprising result from the large IRS micro data set is that, over the entire period 1966–2001, *only the top 10 percent of the income distribution enjoyed a growth rate of total real income (excluding capital gains) equal to or above the average rate of economy-wide productivity growth.* The bottom 90 percent of the income distribution fell behind or even were left out of the productivity gains entirely.

Another way to state our main results is that the *top 1 percent* of the income distribution accounted for 21.6 percent of real total income gains during 1966–2001 and 21.3 percent during the productivity revival of 1997–2001. Still another way is that the top 5 percent of the income distribution earned more of the real 1997–2001 gain in wage and salary income *than the bottom 50 percent.*

Our results show that the dominant share of real income gains accruing to the top 10 percent and the top 1 percent is almost as large for labor income as for total income. This contradicts those economists who believe that growing inequality is entirely a matter of the dominant share of wealth and capital income at the top; for instance, Philip Swagel of the American Enterprise Institute recently stated, "It looks like the gains from the recovery haven't really filtered down. . . . The gains have gone to owners of capital and not to workers."[69] It is not that all the gains went to capital and none to labor; rather, our finding is that most of the gains in labor income, too, went to the very top percentiles.

Many previous papers have documented an increase in American income inequality over the past three decades, but most have used CPS data that, because of top coding, have nothing to say about shifts in the income distribution *within* the top 10 percent of income earners. We document that the top 10 percent of wage and salary earners reaped 45 percent of real income gains during 1966–2001, compared with a 27 percent income share in 1966. Of that 18-percentage-point difference, half is due to an increase of incomes in the 90th percentile and above relative to those below the 90th, whereas the other half is due to increased inequality within the top 10 percent, and especially the gains of the top 0.1 percent compared with the next 9.9 percent.

Our new data on the micro income distribution are accompanied by a review and extension of a large literature in labor economics on inequality. We conclude that there has been virtually no increase in inequality at

69. Quoted in Leonhardt, "U.S. Poverty Rate Was Up Last Year," p. A14.

the bottom as measured by changes in the P50/P10 ratio. The standard explanation, skill-biased technical change, has some role in explaining increased inequality in the group with incomes between the 50th and the 90th percentile, although the literature has raised legitimate questions about the timing of the increase in skill premia and its relation, if any, to the spread of computer technology and the timing of macro productivity growth.

Because our IRS data allow a close look inside the top 10 percent right up to the 99.99th percentile, we emphasize causes of increased inequality at the very top. We combine two separate analyses, the first of the economics of superstars, where technology has broadened audiences and increased the rewards for the very best as compared to the next best, and the second of CEOs, whose current compensation relative to average wages has increased by a large multiple since the 1970s, as has superstar income but perhaps for different reasons. For both analyses we collect new data and cite other papers in concluding that superstars and CEOs account for most of the income earned in the very top reaches of the income distribution.

### *The New Macro Analysis*

The macro analysis of this paper identifies a previously overlooked aspect of inflation and wage dynamics. The most important result is that an acceleration or a deceleration of the productivity growth trend alters the inflation rate by at least one for one in the opposite direction. This is an impact of the *change* in the rate of trend productivity growth and dies out if trend productivity growth stabilizes at a new level, as it did in 1995–2005. Symmetrically, the post-1965 acceleration of inflation was in part caused by the infamous productivity slowdown of those years. Counterfactual simulations of our econometric model suggest that the 1965–80 slowdown in productivity growth boosted inflation on average by about 1.3 percentage points during the 1965–80 simulation period, whereas the 1995–2005 revival of productivity growth held down inflation on average by about 1.2 percentage points over that period.

Linking the macro and the micro analysis, a deceleration of inflation that accompanies a productivity growth revival is good news for wage and salary earners generally. But it does not overturn or in any way conflict with the story of this paper's micro analysis of income distribution.

For a bottom-group wage earner with annual growth in real income after 1995 of 0.5 percent, that real income growth rate would have been –0.5 percent without the productivity growth revival. For a top-group wage earner with annual real income growth of 4.0 percent, the absence of the productivity growth revival would have reduced that growth to 3.0 percent. The effect of faster trend productivity growth after 1995 in reducing the inflation rate was good news for the economy, contributing to improved macroeconomic stability and easing the task of monetary policy. But this productivity-inflation nexus does not alter our main message that increased skewness of the income distribution was responsible for the large divergence between the growth rates of median and mean real wage and salary income.

# *Comments and Discussion*

**Daniel E. Sichel:**[1] Ian Dew-Becker and Robert Gordon have put together a very nice paper that covers some important ground. Although I will highlight some quibbles and questions, my overall view is that the paper is quite interesting and extremely well crafted. The issue the authors examine could be summarized as "Who got the benefits from the increases in labor productivity, in the past decade and over the longer haul?" To address this question, the paper covers four broad issues. First, it reviews the key measurement issues that must be understood before proceeding to a comparison of trends in labor productivity and real compensation per hour. Second, the paper tills some new ground on wage-price dynamics and presents a new inflation equation, an updated version of Gordon's "Goldilocks" equation from several years ago. Third, the paper uses the estimated wage-price dynamics to gauge the effect on the labor share of income when trend productivity growth changes. Finally, the paper turns to a micro analysis of changes in the income distribution, using Internal Revenue Service tax data from 1966 to 2001; importantly, the paper links these data to a measurement framework consistent with the National Income and Product Accounts, allowing the authors to examine which income deciles received gains in real hourly compensation that equaled or exceeded the rate of increase in labor productivity growth, and which received less. I will focus on each of these broad topics in turn.

1. The views expressed are mine alone and should not be attributed to the Board of Governors of the Federal Reserve or other members of its staff.

The first point made in the paper's measurement section is that comparisons of productivity growth and gains in real compensation per hour must be made using comparable data. Some analysts compare productivity in the nonfarm business (NFB) sector with average hourly earnings deflated by the consumer price index and find that productivity has risen significantly faster than real wages. As Dew-Becker and Gordon point out, however, such a comparison can be quite misleading, because the consumer price index used to deflate average hourly earnings differs from the deflator used for NFB productivity, and because the average hourly earnings measure covers the earnings of production and nonsupervisory workers only, not all workers in the sector. Dew-Becker and Gordon steer readers to a more appropriate earnings measure against which to measure changes in labor productivity, namely, a comparably deflated measure of compensation per hour from the Bureau of Labor Statistics' Productivity and Cost (P&C) release. This measure does cover all workers in the sector, and by this measure the difference between the growth rate of labor productivity and that of real compensation per hour over the past several decades is rather small.

Although the authors get to the right numbers for comparing labor productivity and real hourly compensation, their explanation of why P&C compensation per hour is more appropriate for this comparison than average hourly earnings is incomplete. They correctly point out the difference in worker coverage, but they do not mention another very important difference, which is that average hourly earnings excludes benefits whereas P&C compensation per hour includes the value of benefits. The value of output used in the labor productivity figure includes benefits, and it is therefore important to use a compensation measure that also includes benefits, especially given that benefits make up nearly 30 percent of compensation and have risen considerably faster than wages and salaries in recent decades.[2]

The second point emphasized in the paper's measurement section is the well-known fact that the labor share of income has exhibited no discernible

2. According to data from the Employment Cost Index (ECI), benefits made up 29 percent of total compensation in the first quarter of 2005, and their value increased over the period from 1979:4 to 2005:1 at an average annual rate of 5.1 percent, compared with 3.3 percent for wages and salaries. The rate of increase in wages and salaries in the ECI is quite close to that of average hourly earnings over the same period.

trend over the past twenty years or, arguably, the past fifty. Of course, this share could have shifted over these periods if and only if aggregate labor productivity and aggregate real hourly compensation had increased at different rates; thus the stylized fact about constant labor shares implies—and is implied by—the similarity in growth rates of labor productivity and real hourly compensation noted above. I have no quibble with any of this analysis, which is based on the conventional NIPA data. However, this widely accepted result is dependent on the use of those particular data. If one defined investment more broadly than do the NIPA data, a different result about labor shares could emerge. Although this line of argument may seem like a detour from the paper's main story line, it actually could be quite central to the issues raised in the paper.

In particular, the conventional NIPA data fail to capitalize most business investments in intangible assets, but instead count the expense of acquiring or developing these assets as an intermediate purchase.[3] According to two recent papers that I co-wrote with Carol Corrado and Charles Hulten, business investment in intangible capital is about as large as that in tangible capital; moreover, intangible investment has grown considerably more rapidly in recent decades than has tangible business investment, and the stock of intangible capital has grown considerably more rapidly than has the stock of tangible business capital.[4] If these intangible assets are counted as business investment, the share of capital income will show a noticeable upward trend in recent decades, and the share of labor income will show a marked downtrend.[5]

This point has direct implications for the present paper. First, the rapid growth of intangible capital raises questions about the widely accepted stylized fact that the labor share of income has been relatively stable in recent decades. Over periods when the economy is transitioning to a mode of production that is more intensive in intangibles, the shares of labor and

3. Computer software is the most notable exception; other intangible assets that are not currently capitalized in the NIPA data include research and development, other product development costs, investments in firm-specific human capital, and investments in organizational structure (such as Wal-Mart's processes for managing its supply chain).

4. Corrado, Hulten, and Sichel (2005a, 2005b). Closely related work has been done by Nakumura (2003); see also Brynjolfsson (2005) for a review of his work on intangibles.

5. Similarly, results in Corrado, Hulten, and Sichel (2005a, 2005b) imply that, once the additional investments in intangibles are included, labor productivity has increased more rapidly than real hourly compensation in recent decades.

capital could be shifting, even if the NIPA-derived shares are not. Second, to the extent that the returns to intangible capital have accrued to the upper income quantiles (which seems likely, since these returns are likely to accrue to the more educated), the presence of these additional assets could be affecting the upper tiers of the income distribution. Of course, at this point that is only speculation, and more work needs to be done to confirm whether this is so.

The second broad contribution of the Dew-Becker and Gordon paper is the development of a new inflation equation. This equation updates Gordon's "Goldilocks" inflation equation and is similar to it in many ways.[6] Both the new and the old equations include lagged dependent variables entered as four-quarter moving averages for lags one, five, nine, thirteen, seventeen, and twenty-one. Both also include the unemployment gap (measured as the difference between the unemployment rate and a time-varying NAIRU), a term for the relative price of food and energy, a term for the relative price of imports, and terms to capture the Nixon-era price controls. The new specification differs from Goldilocks by including a term for the relative price of medical services, as well as the first and fifth lags of a variable capturing acceleration in labor productivity, measured as the eight-quarter change in the growth of Hodrick-Prescott trend productivity with a smoothing parameter of 6400. In the old equation the productivity terms were entered as the current value and first lag of the deviation of productivity growth from trend.

As the authors demonstrate, the new equation generally passes tests for parameter constancy and performs admirably in out-of-sample forecasting tests.[7] But what should one make of this impressive performance? On the one hand, identifying a single equation that is stable across all of the structural changes in the economy over the past forty years is a remarkable achievement. On the other hand, there is always the risk that the specification just manages to thread the needle of the various specification tests but still is fragile. As is apparent from my description, the new specification has a rather intricate lag structure. Moreover, many practitioners of the dark art of estimating inflation equations know that it is tricky to get a productivity trend to enter an inflation equation significantly. Relatively small

6. Gordon (1998).

7. The one coefficient that is not stable across the sample is that on the term for the relative price of food and energy.

changes in the specification can affect the significance of the productivity terms. Also, the equation includes a large complement of variables representing supply shocks. This choice seems reasonable enough, although it is worth noting that the food, energy, and medical supply shocks in the equation account for about 25 percent of the weight in consumer prices. Pushing this approach to the limit (which Dew-Becker and Gordon do not), one could keep adding still other supply-shock terms, such as a housing-cost supply shock, and eventually get a perfect fit once every component of consumer prices has been included as a supply shock. Given my concern about the potential fragility of the new inflation equation, I would want to see additional sensitivity tests to be completely persuaded that this is a robust specification.

More broadly, even though the parameters of the new inflation specification are stable within its sample period of more than forty years, the specification itself has changed since the Goldilocks version of seven years ago. Thus Gordon's inflation specifications themselves have not been stable across time, even if at particular points in time he has succeeded in developing equations that exhibit within-sample parameter stability. The latest changes to the Gordon inflation equation capture the fact that inflation has been lower over the past ten or so years than the inflation equations of a decade ago might have led one to expect. Dew-Becker and Gordon explain this lower-than-previously-expected inflation both by adding the medical supply-shock term already mentioned and by tweaking the way that trend productivity enters the equation. Others have explained the lower-than-expected inflation through a smaller effect of slack on inflation (that is, a shallower Phillips curve), perhaps linked to a better anchoring of inflation expectations. Although the stability of the new inflation equation is impressive, I do not believe that this paper will win over those who believe that the Phillips curve has become shallower over time.

The third broad area covered by the paper is the dynamics of wages, prices, and productivity. The main contributions here are the development of a wage equation that is combined with a slightly modified price equation to create a system in which wages and prices are dynamically linked. This system is then simulated to gauge how the slowdown in trend productivity around the late 1960s and early 1970s and the speedup in the mid-1990s affected inflation and the labor share of income. The results here seem quite sensible to me, with the most striking finding being the symmetry of the effects of these changes in trend productivity. According to the authors' simulations, the productivity slowdown of the late 1960s and early 1970s boosted inflation by about $1\frac{1}{4}$ percentage points a year, and the mid-1990s speedup held inflation down by about the same amount. Similarly, the productivity slowdown is estimated to have raised the change in the trend labor share of income by 0.2 percentage point a year, whereas the mid-1990s speedup is estimated to have reduced this change by about the same amount.

The final part of the paper turns to micro evidence on the income distribution, using the IRS wage and salary data from 1966 to 2001. As the authors note, other researchers have used these data to look at a similar set of questions.[8] What is new here is the evaluation of gains in real hourly compensation in different slices of the income distribution, relative to the 1.6 percent average annual increase in labor productivity over the same period. The authors' snazzy and quite provocative result is that only those in the 90th percentile and above realized gains in real hourly compensation that matched the rate of productivity growth over the past forty years. Although experts in the literature would undoubtedly find several things to quibble with in this analysis, I suspect that the paper's broad qualitative result would stand up to that sort of scrutiny.

Given that I am not one of those experts, I will make just a few points. First, Dew-Becker and Gordon have data on wages and salaries for the income groups they study, but they do not have data on benefits. They use the aggregate ratio of benefits to compensation to translate their data on wages and salaries into total compensation, thus assuming that benefits affect all quantiles proportionally. One way to evaluate this assumption is with data from the Bureau of Labor Statistics' release on Employer Cost for Employee Compensation (known as the ECEC). These data show, by occupation, the breakdown of nominal compensation into wages and salaries and the various components of benefits. To roughly gauge whether benefits are proportional to overall compensation across income groups, I selected the occupations that had the highest and the lowest overall compensation in 2001 and compared the benefits share of compensation across these occupations.

The highest-paid occupation was "executive, administrative, and managerial" and the lowest-paid was retail trade. In 2001 benefits made up nearly 28 percent of compensation for executives, but just over 20 percent

8. For example, Piketty and Saez (2003).

for workers in retail trade. Moreover, from 1986 to 2001 the benefits share of compensation rose for executives but fell for retail workers.[9] Thus, benefits do not appear to be proportional to income, and this unequal distribution of benefits became more unequal from 1986 to 2001. Although these data represent just an example rather than a comprehensive look at the question, this simple comparison suggests that higher-paid workers had a larger share of benefits in their total compensation than did lower-paid workers and that their share has been rising over time. This example raises the possibility that the present authors' results actually understate the increase in income inequality.

Second, I approve of the fact that the authors avoid identifying a single explanation for the increase in inequality at the very top of the distribution and instead are open to a variety of explanations. I also like their approach of blending explanations—noting, for example, that changes in technology may have fueled both the rise of superstars and the decline in unionization. In addition, I thought it quite clever to track down the data on "The Celebrity 100," professional athletes, and CEOs.

Finally, although I like this part of the paper, it does not fully answer the question of why the income distribution has changed in the way that it has. For example, in their discussion of CEO pay, Dew-Becker and Gordon advocate a variant of the standard managerial power explanation, which they call the "scratch-my-back" model. However, they do not provide a fully satisfying explanation of why the mutual scratching has become more intense in recent decades or of how that explanation fits the facts about the distribution of high-end incomes earlier in the twentieth century. For example, the authors' story implicitly seems to rely on the idea that weak corporate governance in recent years has allowed executives to extract ever-larger compensation packages. However, a recent paper by Carola Frydman and Raven Saks indicates that real gains in compensation for executives were smaller in the 1950s and 1960s, and it seems hard to argue that corporate governance was stronger in that period than it has been more recently.[10]

9. The comparisons start in 1986 because that is the earliest year for which the relevant data are available.

10. Frydman and Saks (2005). Frydman and Saks develop a long and consistent database of executive compensation. They also provide evidence that marginal tax rates affect the relationship between stock market valuation and executive pay. In particular, Frydman

To sum up, this paper makes several valuable contributions to the debate over who has benefited most from recent productivity gains. The conclusions are quite sobering and raise an interesting set of questions about the linkages among labor productivity growth, compensation, and well-being.

**Robert Topel:** According to the Bureau of Labor Statistics, the highest-earning profession in the United States is—ready?—*economics teachers.* Of course, this is as it should be. Economic forces have increased the demand for the particular skills we teach, which has raised the price for our services. Our wages have grown faster than aggregate productivity over the past thirty years. My back-of-the-envelope calculation indicates that the real wage of starting assistant professors in top university departments of economics roughly tripled between 1979 and 2000.

As its title suggests, this paper by Ian Dew-Becker and Robert Gordon seeks to understand who has gained from increased productivity in the United States. It is motivated by three well-known facts about wage and productivity growth and the distribution of income:

—Labor's share of national income in the United States has been remarkably stable in the long run, averaging about 70 percent. As a corollary, long-run growth in both productivity and compensation per worker has been roughly equal.

—Second, despite the correspondence between compensation and productivity growth over long periods, labor's share may fluctuate over shorter periods. Periods of accelerated productivity growth, as occurred in the late 1990s, may reduce labor's share.

—Third, even over the long run—say, from 1970 to today—the distribution of gains from productivity growth has been highly uneven. Although the gains from rising productivity accrue to labor, by conventional measures the wage and salary income of the median worker has barely increased, if at all. Only in the upper reaches of the wage distribution have wages grown as much as productivity, reflecting a substantial increase in wage inequality in the United States.

---

and Saks' empirical work indicates that increases in the stock market value of firms lead to smaller increases in the real before-tax value of executive compensation when marginal tax rates are higher than when they are lower.

Dew-Becker and Gordon document and seek to explain these features of the data, offering two distinct and (in my view) noncomplementary analyses. The first combines careful measurement of relevant series from the NIPA data with a time-series econometric model of NIPA aggregates—average wages, productivity, the price level, and so on. Beyond the niceties of such models, which I shall leave to the cognoscenti, their contribution is to explain short-term fluctuations in labor's share in response to changes in productivity growth. I argue that the framework ignores some of the fundamental economic forces at work: the long-run stability of labor's share, combined with short-run changes in that share when productivity growth changes, is implied by the structure of modern growth models.

The second is a detailed analysis of individual tax returns, based on IRS data, that attempts to link macro-level productivity growth to the distribution of income. The point here is to document who gained from rising productivity and to resolve the "puzzle" of rising average productivity but stagnant median earnings. The resolution lies in the difference between an average and a median. Although this is the correct answer, I argue that this was not really a puzzle to knowledgeable observers of income distribution. Along the way, Dew-Becker and Gordon offer a critique of the "skill-biased technical change" (SBTC) explanation for rising inequality in the United States, and their own interpretation of what the causes might be. I did not find this analysis compelling.

DETERMINANTS OF LABOR'S SHARE IN THE LONG AND THE SHORT RUN. The "macro" section of the paper makes two contributions. The first is a careful measurement exercise applied to the sources of national income. I thought this was quite useful. The second is a time-series model of wage and price dynamics, which then "determines" labor's share. Key to this is that the model has enough degrees of freedom that labor's share may rise or fall in response to a change in productivity growth, depending on the temporal patterns with which $w$ and $p$ adjust. I have no quarrel with the details of this exercise, but I have two critical comments about its place in this paper. First, it is entirely disconnected from what I see as the paper's focus, which is to document the *distribution* of income gains from recent productivity growth. The paper's central message would not change if this model were not in it. Second, it bears noting that observed patterns of adjustment in income shares are consistent with certain economic fundamentals that the paper ignores. A macroeconomic model of inflation and wage dynamics is not necessary to understand why labor's share may fluctuate about a long-run stable value; real (that is, nonmonetary) factors will do the trick.

The observation that labor's share is fairly constant in the long run is one of the foundations of modern theories of economic growth, beginning with Nicholas Kaldor's observation that the capital-output ratio is constant and Robert Solow's contribution to growth accounting in the 1950s. Thus Solow, Hirofumi Uzawa, and Robert Lucas,[1] among others, posit that national output is produced with constant returns and labor-augmenting technical progress:

$$Y = F\left(AL, K\right). \tag{1}$$

Here $A$ is the state of labor-augmenting technical progress. Then $K/L$ grows at the same rate as $A$, yielding constant income shares for labor and capital. In this framework Paul Krugman's comment that "the essential arithmetic says that *long-term* growth in living standards . . . depends almost entirely on productivity growth" is dead on.[2]

In this framework, changes in labor's share are determined by the elasticities of supply of labor and capital and the elasticity of substitution between them. Assume that labor is in fixed supply, and denote the supply elasticity of capital by $\eta$. Let $\sigma$ be the elasticity of substitution between labor and capital. Labor's share is the ratio of the marginal and average products of labor, which responds to the change in productivity:

$$\frac{d \ln\left(MP_L / AP_L\right)}{d \ln A} = \left(\sigma - 1\right) s_k \left(\frac{1}{\sigma + s_k \eta}\right), \tag{2}$$

where $s_k$ is capital's share. According to equation 2, wages will grow more slowly than productivity if $\sigma < 1$ and the supply of capital is not perfectly elastic. Both of these conditions describe the "short run": the ease of substitution between capital and labor surely rises with the length of the adjustment period, and the evidence is that capital is in (very) elastic supply in the long run. Thus standard models of economic growth, which are the foundation for the connection between productivity growth and rising living

1. Solow (1956); Uzawa (1965); Lucas (1988).

2. Emphasis added. Krugman is one of four *New York Times* columnists quoted in the paper.

**Figure 1. Wages at Selected Percentiles and Productivity, 1963–2000**

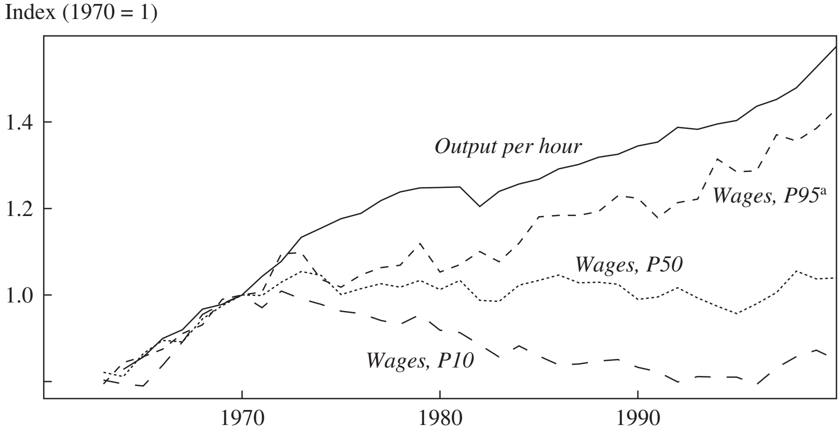


Sources: National Income and Product Accounts, Bureau of Economic Analysis; author's calculations.
a. "Pn" refers to the nth percentile.

standards mentioned by the authors, will do the trick. Short-run changes in labor's share in response to changes in productivity are no mystery.

PRODUCTIVITY GROWTH AND RISING INEQUALITY. The remainder of the paper offers an accounting of income growth and its distribution, based on IRS tax return data, and an interpretation of the data. The motivation for this is illustrated in figure 1 above, which shows indices for labor productivity and weekly wages at the 95th, 50th, and 10th percentiles of the distribution of wages of full-time, full-year male workers. The top line represents productivity, which grew by about 60 percent between 1970 and 2000. The next line down is the wage of men at the 95th percentile, measured from the March Current Population Survey (CPS), which grew by slightly over 40 percent, and so on. The relative wage indices clearly illustrate the trend toward rising inequality in the United States. As important for this paper, wages measured in the CPS grew more slowly than productivity even for men at the 95th percentile of the wage distribution. Real wages for the median male worker showed essentially no long-term growth. These facts are well known.

There are basically two ways to reconcile the wage and productivity data—to make things add up, as it were. First, total compensation (the relevant concept for NIPA) may rise faster than wages because nonwage benefits—which are not counted in the CPS or, for the most part, in IRS data—become a more important component of compensation. The second is that the financial gains from rising productivity have accrued to very high earners, while the vast majority of wage earners have gained proportionally less. There is some truth to both explanations: it is widely known that benefits are a rising portion of compensation, and rising inequality is perhaps the most widely studied empirical fact in economics over the past twenty years.

Although they acknowledge that nonwage compensation has grown more rapidly, on average, than wages, Dew-Becker and Gordon dismiss the importance of this fact for understanding the level and distribution of gains from rising productivity. Indeed, they assume, without discernable foundation, that relative growth of nonwage compensation affects all percentiles of the income distribution equally. Thus they "apply the decline in the NIPA wage share of compensation equally to all percentiles," which means that compensation relatives are unaffected by their adjustment, and that compensation growth for the median earner is not adjusted upward by much. But that is not what happened. Much of nonwage compensation comes in the form of a fixed cost per employee—the rising cost of health care benefits is a good example—that have a larger impact on the median earner than on high earners. This issue was studied by Brooks Pierce,[3] who used internal BLS data on the employment cost index (ECI) to estimate the evolution of "compensation inequality." Pierce finds that the fraction of compensation taken as wages is lower at the 10th percentile than at the 50th, and higher at the 50th than at the 90th.

These findings reflect, in part, the income elasticity of demand for benefits. If the real incomes of very low wage workers have declined, as the data clearly indicate, the first thing to go may be benefits, especially if the costs of those benefits have risen. Thus the fraction of low-wage workers with employer-provided health insurance has fallen dramatically, reflecting powerful substitution and income effects on demand. Pierce also estimates that income elasticities decline monotonically with income. Overall, the data indicate that compensation growth exceeded wage growth for the middle half of the wage distribution, but not in the extremes. In the left tail of the distribution, compensation grew more slowly than wages—exacerbating inequality—while compensation and wages grew at roughly

3. Pierce (1999).

the same rate above the 80th percentile. Pierce's data have little to say about the extreme extremes of the wage distribution, but it is plausible that wage and salary income is the whole story at the top, where the data suggest that the income elasticity of demand for benefits would be well below 1.0.

The upshot is that more of productivity growth found its way to the typical worker than Dew-Becker and Gordon indicate—the connection between "living standards" and productivity is not so tenuous as they would have us believe, although Pierce's findings are not so strong as to negate the distributional issues studied here. Further, Pierce's results on the distribution of benefits reinforce the conclusion that the fortunes of low-skill workers declined after 1973. Surprisingly, Dew-Becker and Gordon deny the latter point. They state that relative wage inequality below the median of the wage distribution has not changed over time, and they cite a recent paper by Autor, Katz, and Kearney as reaching a similar conclusion.[4] As is widely known, and indeed *shown* by Autor, Katz, and Kearney, the wage differential, based on survey data, between the 50th and the 10th percentile widened through the mid-1980s—the real wages of low-skill workers were declining—after which it has been roughly stable. Inequality in the upper half of the distribution continued to rise.

The IRS data employed by Dew-Becker and Gordon are extremely useful for demonstrating that rising inequality was itself skewed toward the upper reaches of the wage distribution. In survey data like the CPS, earnings responses are "top coded" to protect confidentiality, and so the data are not very informative about the earnings of workers above about the 95th percentile. Although the IRS data have many disadvantages, they are valuable because they record wage and salary income for individuals earning, say, $10 million a year. Thus they provide detail on what has happened at very high altitudes. In a nutshell, the authors find that wage and salary growth has been highly skewed toward the very top, confirming the findings of Thomas Piketty and Emmanuel Saez,[5] who also used IRS data.

Why did inequality rise? Why did the gains from rising productivity find their way mainly to the top of the wage distribution? Dew-Becker and Gordon eschew an explanation based on their characterization of SBTC, preferring a combination of "superstars" and CEOs engaged in something

4. Autor, Katz, and Kearney (2005).
5. Piketty and Saez (2003).

called the "scratch-my-back" model, where executives determine each others' pay.

SKILL-BIASED TECHNICAL CHANGE. The marginal productivity theory of income distribution puts some fairly severe restrictions on where one might look for explanations of rising wage inequality. Consider just two worker types, skilled ($S$) and unskilled ($U$). The ratio of their wages ($W_j$) is determined by the ratio of the value of their marginal products:

$$\frac{W_S}{W_U} = \frac{P_S}{P_U}\frac{MP_S}{MP_U}.$$

Here $P_j$ is a price index for goods made by skilled or unskilled labor, and $MP_j$ is the marginal productivity of a worker of type $j$. So inequality will increase if either the ratio of prices or the ratio of productivities rises. The former can be affected by, say, increased trade with less developed countries, and the latter by changes in factor ratios (due to immigration, investment in human capital, and so on) or by technical progress that raises the productivity of one group relative to the other. An example of the last effect is SBTC, which Dew-Becker and Gordon describe as a "shift in . . . demand toward more-skilled workers."

To Dew-Becker and Gordon, however, SBTC is evidently a much narrower concept identified with the increase in computing power and applications: "If SBTC had been a major source of the rise in inequality, we should have observed an increase in the relative wages of those most directly skilled in the development and use of computers." This is wrong. The effect of technical progress on the wage of a particular type of worker—say, a programmer or a developer—depends on the elasticity of demand for the product, on the elasticity of supply of individuals with the requisite skills for the task, and on the amount of skills required. Even if technical progress has simply raised the productivity of a given quantity of computer jocks, the wage of computer jocks may easily fall.[6]

Dew-Becker and Gordon also criticize SBTC because wage inequality narrowed in the United States from the 1940s to 1970. "Are we to believe that technical change over 1920–70 was 'unskilled biased'?" they ask.

6. A simple example: Assume constant returns to scale in strawberry harvesting. Assume that technical progress raises the productivity of strawberry pickers by 1 percent, which will reduce the price of strawberries by 1 percent as well. Then the demand for strawberry pickers will increase if the elasticity of demand for strawberries exceeds 1.0.

Well, actually it could have been. And it is salient that inequality of measured skills declined over this period, as greater proportions of young people graduated from high school and college, while agriculture declined as a source of employment. This has a direct effect on inequality—there are more workers with skills—as well as an indirect effect through factor proportions and hence prices.[7] Exactly this type of phenomenon occurred in the 1970s, when the labor supply effect of college-educated baby-boom cohorts caused the college–high school wage ratio to decline, even as overall wage inequality was rising. The authors dismiss these effects: rising relative wages of less skilled workers "depended perhaps more on the strength of unions, and the relative absence of immigration and imports." Although these may be important, the authors present not an iota of evidence in support of this sentiment.

More-direct evidence of a role for SBTC comes from the relative employment growth rates of industries with different skill intensities. Dividing industries into quartiles by years of average schooling of the workforce, Kevin M. Murphy and Finis Welch find that these growth rates are rank ordered: from 1970 to 1998 the most skill intensive industries saw their employment grow the most and the most steadily—industries in the top quartile expanded their employment by about 75 percent—while employment in the least skill intensive industries steadily contracted.[8] Evidently the demand for skilled people increased—as is reflected in the prices they commanded in the labor market.

Dew-Becker and Gordon's last critique of SBTC is that it "fails to explain the absence of an increase in income inequality in Europe despite the free flow of technology across borders." This "fact" is news to me. From my own work in equality-obsessed Sweden, I know that wage and income inequality increased there, as well as in many other European countries.[9] For a broader perspective I consulted the data of the Organization for Economic Cooperation and Development (OECD), which calculates income inequality from household surveys in its member countries. With the exception of the Netherlands, the OECD data show an increase in income inequality, as measured by Gini coefficients, in fourteen of fif-

7. See, for example, Topel (1997).
8. Murphy and Welch (2001).
9. Edin and Topel (1997); Fredrikson and Topel (2005).

teen member countries between the mid-1980s and 2000.[10] It is true that inequality is *higher* in the United States than in other OECD countries, and it is probably true that European labor market institutions have constrained the increase in observed wage inequality that would otherwise have occurred, creating other social costs and distortions in the process. But inequality in Europe has risen nevertheless.

From my reading of the evidence, SBTC still looks good as *an* (perhaps not *the*) explanation for rising wage inequality. It is probably the most important explanation, but rankings need not concern us. Having dismissed it altogether, however, and having oddly concluded that inequality did not increase among low-wage workers, Dew-Becker and Gordon offer two alternative explanations for widening earnings differentials among very high earners. The first is the phenomenon of superstars—the idea that the output of a single individual can now reach more users than before. The second is the scratch-my-back model of CEO pay.

FACTORS AFFECTING PAY AT THE TOP OF THE WAGE DISTRIBUTION. Consider superstars, which the authors identify with professional athletes and entertainers. Because of advances in communications technology, it is now possible for many more users to enjoy the output of the "best" producer: millions could watch Michael Jordan without going to a Chicago Bulls game. This is surely an aspect of SBTC: the "span of influence"—to use one of Sherwin Rosen's favorite phrases—has been increased by technological change, raising the relative productivity of a talented few. To Dew-Becker and Gordon this is an understandable "equilibrium" phenomenon when the subjects are Michael Jordan, Oprah Winfrey, or (dare I say) the writers of economics textbooks.[11] But those examples account for only a small portion of the increase in income among the very rich. CEOs, it is alleged, have grabbed most of the rest. And since the authors do not understand why their earnings have increased, it must be bad.

One might first quibble with the facts. Although it is true that a small cadre of CEOs have experienced astonishing increases in pay, I think the authors fail to establish that they account for a large portion of the

10. Förster and d'Ercole (2005). See also Smeeding (2002).

11. The authors repeatedly cite pay differences between the United States and Europe. For example, they note that CEOs in Europe and other countries earn less than in the United States. They might have also noted that European athletes and entertainers earn less than those in the United States.

extreme tail of the distribution: more than half of earnings increases above the 99.99th percentile, as they claim. They cite evidence that the top five executives in 1,500 firms received average pay in 2001 of $6.4 million: "This is 7,500 people making $6.4 million each." No, it isn't. It is 7,500 people with an *average* income of $6.4 million, many of whom would not qualify for the 99.99th percentile of the wage distribution.[12] *This is precisely the distinction that motivated this paper,* and it is misapplied here.

Much of the literature on CEO pay is of the widely used "$y = \beta x$ + my theory" variety. That is, failing to explain some phenomenon with simple economics, all constraints are off. This paper is no exception. The scratch-my-back model avers that CEOs support each others' exorbitant pay increases, and things just go wild. Ironically, the authors support this notion by quoting at length yet another *New York Times* journalist who happens to hold a similar view.

It is popular to conclude that top executives are overpaid, which must mean that investors could get the same performance for a fraction of the costs, but do not. Maybe. Other authors, such as Michael Jensen and Kevin J. Murphy, have argued that CEOs are (or at least were) underpaid. Maybe. So far as I can tell, nothing in this literature provides convincing evidence either way. The scratch-my-back discussion does not advance the ball.

Against these arguments it must loom large that the expansion of pay at the very top of the wage distribution is a magnification of what has happened at lower altitudes, and a reflection of what has happened to superstars. From 1970 to 2000 wage growth increased smoothly and monotonically across percentiles of the wage distribution, from the bottom up to the 95th percentile. Why did the earnings of college graduates rise relative to those of high school graduates? I think we have a good idea. Why did the earnings of those with advanced degrees rise relative to the earnings of college graduates? I think we have a good idea. Why did the earnings of economics teachers rise relative to those of workers in other fields? I think we have a good idea. Why did the earnings of superstars rise so much? I think we have a good idea. Why did the rich get richer? Why is it something else?

12. The *Forbes* survey of executive compensation reports the pay of the 800 highest-compensated CEOs in the United States. The median total compensation is about $2.5 million, which according to the authors is outside the 99.99th percentile.

13. Jensen and Murphy (1990).

**General discussion:** Edmund Phelps was skeptical of the paper's link between U.S. productivity growth and labor's share of income. The United States is, after all, only part of the world economy, currently accounting for about a quarter of gross world product. Developments abroad could help account for variations in capital income, thus calling the tight linkage into question. In the two main recent periods of rapid U.S. productivity growth, productivity rose even more in some other parts of the world: in continental Western Europe and Japan in the 1960s, and in Eastern Europe and China in the 1990s. These developments affected world real interest rates and the growth of income from wealth relative to the stock of U.S. business assets, and so provide an alternative explanation for variations in the nonlabor components of U.S. income.

Addressing Robert Topel's discussion of the paper, Phelps noted that many models of economic activity do not fix the relative shares of income to labor and capital in the way that the familiar Cobb-Douglas production function does. Instead there can be several different factors of production and business assets, such as firm-specific, trained employees, which can lead to changes in labor's share both in the short and in the long run. Richard Cooper challenged the authors' claim that, because income inequality has not increased in Europe and some other countries, skill-biased technical change lacks empirical support outside the United States. He pointed out that Gini coefficients in Mexico, India, China, Russia, and some other emerging economies all show an increase in inequality. Indeed, the world Gini coefficient has moved toward greater inequality because China and India, with their huge populations, have also been growing more rapidly than the rich economies.

Lawrence Mishel reasoned that what matters for wage inequality is not whether technological change affects the need for skills but whether that impact has accelerated: Has the expanded demand for skills grown faster than their supply? He observed that the supply of skills and education has grown rapidly: in the past thirty years the share of the workforce with a college degree has doubled, while the share who are high school dropouts has fallen to a third of what it was. James Duesenberry noted that particular developments affecting incomes can be identified at both extremes of the income distribution. Immigration and the decline of unions have held down wages at the low end, while intense competition for high-quality professionals has lifted incomes at the very high end. Most of the increases at the top involve occupations, such as medicine, law, finance, and corporate management, that require more than a basic college education. And there has been increased competition for people who are exceptionally successful. On the supply side, the most highly rated professional schools have not increased their output at the same rate as demand has grown, and so their graduates command a higher income premium than they used to.

David Laibson suggested that some of the rise in U.S. inequality is driven by a change over the past century in the way men and women choose mates. Many of today's young adults marry people who are similar intellectually and have the same ambitions and motivations. This assortive matching not only creates households in which both spouses are equally highly paid, but also tends to produce offspring with more extreme attributes, not only in terms of IQ but also in terms of motivation and training.

Responding to the discussion, Dew-Becker and Gordon pointed out that the two discussants disagreed about the distributional pattern of employee benefits. Sichel had offered evidence that benefits had risen at the top and declined at the bottom, implying that the paper understated the increase of inequality. In contrast, Topel had cited evidence indicating that benefits matter most in the middle of the distribution and less at either tail and that, above the 80th percentile, benefits and wages grew at roughly the same rate, implying that the paper did not misstate the increase in inequality for the top 20 percent by omitting separate data on benefits.

## References


Abraham, Katherine G., James R. Speltzer, and Jay C. Stewart. 1999. "Why Do Different Wage Series Tell Different Stories?" *American Economic Review* 89, no. 2: 34–39.

Acemoglu, Daron. 2002. "Cross-Country Inequality Trends." Working Paper 8832. Cambridge, Mass.: National Bureau of Economic Research (March).

Atkeson, Andrew, and Lee E. Ohanian. 2001. "Are Phillips Curves Useful for Forecasting Inflation?" Federal Reserve Bank of Minneapolis *Quarterly Review* 25, no. 1: 2–11.

Autor, David H., Lawrence F. Katz, and Melissa S. Kearney. 2005. "Trends in U.S. Wage Inequality: Re-assessing the Revisionists." Working Paper 11627. Cambridge, Mass.: National Bureau of Economic Research (September).

Ball, Laurence, and Robert Moffitt. 2001. "Productivity Growth and the Phillips Curve." 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.

Beaudry, Paul, and David A. Green. 2002. "Changes in U.S. Wages 1976–2000: Ongoing Skill Bias or Major Technological Change?" Working Paper 8787. Cambridge, Mass.: National Bureau of Economic Research (February).

Bebchuk, Lucian, and Yaniv Grinstein. 2005. "The Growth of Executive Pay." Working Paper 11443. Cambridge, Mass.: National Bureau of Economic Research (June).

Bradbury, Katherine, and Jane Katz. 2002. "Are Lifetime Incomes Growing More Unequal? Looking at New Evidence on Family Income Mobility." Federal Reserve Bank of Boston *Regional Review* 12, no. 4: 3–5.

Brynjolfsson, Erik. 2005. "Remarks." In *Measuring Capital in the New Economy,* edited by Carol Corrado, John Haltiwanger, and Daniel Sichel. National Bureau of Economic Research and University of Chicago Press.

Card, David, and John E. DiNardo. 2002. "Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles." *Journal of Labor Economics* 20, no. 4: 733–83.

Corrado, Carol, Charles Hulten, and Daniel Sichel. 2005a. "Measuring Capital and Technology: An Expanded Frameword." In *Measuring Capital in the New Economy,* edited by Carol Corrado, John Haltiwanger, and Daniel Sichel. National Bureau of Economic Research and University of Chicago Press.

__________. 2005b. "Intangible Capital and Economic Growth." Board of Governors of the Federal Reserve (October).

Edin, Per-Anders, and Robert Topel. 1997. "Wage Policy and Restructuring: The Swedish Labor Market Since 1960." In *The Welfare State in Transition: Reforming the Swedish Model*, edited by Richard B. Freeman, Robert Topel, and Birgitta Swedenborg. University of Chicago Press for the National Bureau of Economic Research.

Förster, Michael, and Marco Mira d'Ercole. 2005. "Income Distribution and Poverty in OECD Countries in the Second Half of the 1990s." OECD Social Employment and Migration Working Paper 22. Paris.

Fredrikson, Peter, and Robert Topel. 2005. "Wage Determination and Employment in Sweden Since the Early 1990s." Paper prepared for *The Swedish Welfare State Revisited—What Have We Learned?* (October).

Frydman, Carola, and Raven E. Saks. 2005. "Historical Trends in Executive Compensation: 1936–2003." Harvard University (November).

Galí, Jordi, and Mark Gertler. 1999. "Inflation Dynamics: A Structural Econometric Analysis." *Journal of Monetary Economics* 44, no. 2: 195–222.

Galí, Jordi, Mark Gertler, and David Lopez-Salido. 2001. "European Inflation Dynamics." *European Economic Review* 45, no. 7: 1237–70.

Goldin, Claudia, and Robert A. Margo. 1992. "The Great Compression: The Wage Structure in the United States at Mid-Century." *Quarterly Journal of Economics* 107 (February): 1–34.

Gordon, Robert J. 1975. "Alternative Responses of Policy to External Supply Shocks." *BPEA,* no. 1: 183–206.

_________. 1977. "Can the Inflation of the 1970s Be Explained?" *BPEA,* no. 1: 253–77.

_________. 1982. "Inflation, Flexible Exchange Rates, and the Natural Rate of Unemployment." In *Workers, Jobs, and Inflation*, edited by Martin N. Baily. Brookings.

_________. 1998. "Foundations of the Goldilocks Economy: Supply Shocks and the Time-Varying NAIRU." *BPEA,* no. 2: 297–333.

_________. 2003. "Exploding Productivity Growth: Context, Causes, and Implications." *BPEA,* no. 2: 207–79.

Gordon, Robert, and Stephen R. King. 1982. "The Output Cost of Disinflation in Traditional and Vector Autoregressive Models." *BPEA,* no. 1: 205–42.

Gottschalk, Peter, and Sheldon Danziger. 2005. "Inequality of Wage Rates, Earnings, and Family Income in the United States, 1975–2002." *Review of Income and Wealth* 51, no. 2: 231–54.

Hubbard, R. Glenn. Forthcoming. "Pay without Performance: A Market Equilibrium Critique." *Journal of Corporation Law*.

Jensen, Michael C., and Kevin J. Murphy. 1990. "Performance Pay and Top Management Incentives." *Journal of Political Economy* 98, no. 2: 225–64.

Juhn, Chinhui, Kevin M. Murphy, and Robert H. Topel. 2002. "Current Unemployment, Historically Contemplated." *BPEA,* no. 1: 79–116.

Katz, Lawrence F., and Kevin M. Murphy. 1992. "Changes in Relative Wages, 1963–1987: Supply and Demand Factors." *Quarterly Journal of Economics* 107 (February): 35–78.

Kopczuk, Wojciech, and Emmanuel Saez. 2004. "Top Wealth Shares in the United States, 1916–2000: Evidence from Estate Tax Returns." Working Paper 10399. Cambridge, Mass.: National Bureau of Economic Research (March).

Krugman, Paul. 1990. *The Age of Diminished Expectations: U.S. Economic Policy in the 1990s.* MIT Press.

Ledbetter, Mark A. 2004. “Comparison of BEA Estimates of Personal Income and IRS Estimates of Adjusted Gross Income.” *Survey of Current Business* 84 (November): 9–14.

Lucas, Robert E. 1988. “On the Mechanics of Economic Development.” *Journal of Monetary Economics* 22, no. 1: 3–42.

Mishel, Lawrence, Jared Bernstein, and Sylvia Allegretto. 2005. *The State of Working America 2004/2005.* Cornell University Press.

Mischel, Lawrence, Jared Bernstein, and John Schmitt. 2001. *The State of Working America 2000/2001.* Cornell University Press.

Murphy, Kevin M., and Finis Welch. 2001. “Industrial Change and the Demand for Skill.” In *The Causes and Consequences of Increasing Inequality,* edited by Finis Welch. University of Chicago Press.

Nakamura, Leonard. 2003. “A Trillion Dollars a Year in Intangible Investment and the New Economy.” In *Intangible Assets,* edited by John R. M. Hand and Baruch Lev. Oxford University Press.

Park, Thae S. 2002. “Comparison of BEA Estimates of Personal Income and IRS Estimates of Adjusted Gross Income: New Estimates for 2000 and Revised Estimates for 1999.” *Survey of Current Business* 82 (November): 13–20.

Pierce, Brooks. 1999. “Compensation Inequality.” Working Paper 323. U.S. Department of Labor, Bureau of Labor Statistics.

Piketty, Thomas, and Emmanuel Saez. 2003. “Income Inequality in the United States, 1913–1998.” *Quarterly Journal of Economics* 118: 1–39.

Rosen, Sherwin. 1981. “The Economics of Superstars.” *American Economic Review* 71, no. 5: 845–58.

Samuelson, Paul A., and Robert M. Solow. 1960. “Analytical Aspects of Anti-Inflation Policy.” *American Economic Review Papers and Proceedings* 50, no. 2: 177–94.

Sims, Christopher A. 1987. “Comment.” *BPEA,* no 1: 117–20.

Smeeding, Timothy M. 2002. “Globalization, Inequality and the Rich Countries of the G-20: Evidence from the Luxembourg Income Study.” Presented at the G-20 Meeting on Globalization, Living Standards and Inequality, Sydney, May 26–28.

Solow, Robert. 1956. “A Contribution to the Theory of Economic Growth.” *Quarterly Journal of Economics* 70, no. 1: 65–94.

Staiger, Douglas, James H. Stock, and Mark W. Watson. 1997. “The NAIRU, Unemployment, and Monetary Policy.” *Journal of Economic Perspectives* 11, no. 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.

Stock, James H., and Mark W. Watson. 1996. "Asymptotically Median Unbiased Estimation of Coefficient Variance in a Time Varying Parameter Model." Technical Working Paper 201. Cambridge, Mass.: National Bureau of Economic Research.

__________. 2002. "Has the Business Cycle Changed and Why?" In *NBER Macroeconomics Annual 2002.* MIT Press.

__________. 2003. "Has the Business Cycle Changed?" In *Monetary Policy and Uncertainty: Adapting to a Changing Economy.* Federal Reserve Bank of Kansas City.

Tetlow, Robert J., and Brian Ironside. 2005. "Real-Time Model Uncertainty in the United States: The Fed from 1996–2003." Working paper. Board of Governors of the Federal Reserve (February).

Topel, Robert H. 1997. "Factor Proportions and Relative Wages: The Supply-Side Determinants of Wage Inequality." *Journal of Economic Perspectives* 11, no. 2: 55–74.

Uzawa, Hirofumi. 1965. "Optimal Technical Change in an Aggregative Model of Economic Growth." *International Economic Review* 6, no. 1: 18–31.