

A Bit of a Miracle No More: The Decline of the Labor Share

BY ROC ARMENTER

How is income divided between labor and capital? Every dollar of income earned by U.S. households can be classified as either labor earnings — wages and other forms of compensation — or capital earnings — interest or dividend payments and rent. The split between labor and capital income informs economists' thinking on several topics and plays a key role in debates regarding income inequality and long-run economic growth. Unfortunately, distinguishing between labor and capital income is not always an easy task.

Until recently, the division between labor and capital income had not received much attention. The reason was quite simple: Labor's share never ventured far from 62 percent of total U.S. income for almost 50 years — through expansions, recessions, high and low inflation, and the long transition from an economy primarily based on manufacturing to one mainly centered on services. As it happened, the overall labor share remained stable as large forces pulling it in opposite directions canceled each other out — a coincidence that John Maynard Keynes famously called “a bit of a miracle.” But the new millennium marked a turning point: Labor's share began a pronounced fall that continues today.

Why did the labor share lose its “miraculous” stability and embark on a steep decline? To investigate this shift, economists must first be sure they are measuring the labor share correctly. Could measurement problems distort our understanding of what has happened to the labor share over time? In this article, I explain the inherent challenges in measuring the labor share and introduce several alternative definitions designed to address some of the measurement problems. As we will see, the overall trend is confirmed across a wide range of definitions.

Economists do not yet have a full understanding of the causes behind the labor share's decline. We can make some progress, though, by noting the impact of wage and produc-

tivity trends and shifts between industries. Finally, I discuss several popular hypotheses, based on concurrent phenomena, such as widening wage inequality and globalization, that may account for the labor share's sharp decline.

MEASURING THE U.S. LABOR SHARE

By construction, all income accounted for in the U.S. economy must be earned either by capital or labor.¹ In some cases, we can easily see whether our income comes from labor or capital: when we earn a wage or a bonus through our labor or when we earn interest from our savings or investment account, which is attributed to capital income, despite the fact that most of us would not think of ourselves as investors. However, it is not always immediately apparent that *all* income eventually accrues to either capital or labor. For example, when we buy our groceries — creating income for the grocer — we are only vaguely aware that we are also paying the producers, farm workers, and transporters as well as for the harvesters, trucks, trains, coolers, and other capital equipment involved in producing and distributing what we purchase. However, when the Bureau of Economic Analysis (BEA) constructs the national income and product accounts, it combines data from expenditures and income to ensure that every dollar spent is also counted as a dollar earned by either capital or labor.

Of course, nothing is ever so simple when it comes to economic statistics. First, we lack the detail necessary to split some components of the income data between labor and capital returns. As I will show, the foremost example is the income of self-employed workers, who simply collect the income of

Roc Armenter is a vice president and economist at the Federal Reserve Bank of Philadelphia. The views expressed in this article are not necessarily those of the Federal Reserve.

their business without distinguishing whether it resulted from their work or their investment. In addition, the housing and the government sectors have their total income arbitrarily assigned as labor and capital income, respectively, in the national income accounts. The Bureau of Labor Statistics (BLS) publishes the most widely used estimate of the labor share, which combines several data sources and estimates in order to get around some of the measurement problems. However, these problems remain significant enough that economists routinely create alternative definitions and compare results across them, since a single definition of the labor share is unlikely to fit all purposes.

And what would these purposes be? First, workers and investors respond differently to the same economic conditions and policies. So if we wish to understand how aggregate output will respond, we need to know how to weigh the responses of workers and investors. Second, assets are notoriously unevenly distributed across households; hence, an increase in the share of income earned by capital contributes to income inequality, as richer households would receive an even larger share of total income. Third, the tax code treats labor and capital income differently. Labor income is subject to payroll taxes and the usual income tax rate schedule. Corporate profits (the main source of capital income) are subject to corporate taxes as well as dividends and capital gains taxes when profits are distributed to households. A shift in the labor share will impact not only tax revenues but also how the burden of taxation is distributed across households. Economists also need an estimate of the labor share when determining how much of economic growth can be attributed to labor force growth, capital accumulation, or technological changes — which in turn are key inputs for long-term growth forecasts.

Components of income. The BEA measures output, or gross value added, two different ways in its national income and product accounts (NIPA).² The *expenditure approach* aims to measure the total amount spent on goods and services throughout a year; the *income approach* instead adds up all the income earned by households. In theory, both measures should yield the same number. In practice, alas, they do not. The discrepancy is due to data limitations and measurement error, though the discrepancy is quite small.

The labor share is measured using the income approach. Every dollar of output must be earned by *factors of production* and distributed to households. What exactly is a factor of production? Were we to measure the output of a factory, we would count as factors of production the workers and managers, all the equipment, the building and land

occupied by the factory as well as the electricity, security service, and all the other intermediate inputs used. But because we are measuring the output of the whole economy, we must recognize that the intermediate goods, utilities, and services were produced by some other firm, which in turn uses its own factors of production. Were we to check with, say, the firm producing electricity, we would once again find some workers and managers, equipment, and so on. Now, we could try to track each and every input of production in the U.S., but we would quickly realize that the only factors of production whose income accrues directly to households are labor and capital.

The BEA classifies output into seven groups, as detailed in the table. The second column provides the share of each component relative to the total for 2013. The classification of most income sources as capital or labor income is quite unambiguous. For example, compensation of employees clearly accrues to labor, while corporate profits, rental income, and net interest income are returns to capital. Of the three remaining components, the main challenge is proprietor's income.³

Proprietor's income is defined as the income of sole proprietorships and partnerships — in other words, the income of self-employed individuals.⁴ There is no question that their income is the result of both labor and capital. For example, a freelance journalist may work long hours to document and write a story using a computer and a camera that she or he financed through savings. However, self-employed individuals have no need, economic or fiscal, to distinguish between wages and profits. However, economists do.

The main BLS measure. The BLS is well aware of these problems and goes to great lengths to disentangle proprietor's income into its labor and capital income components. First, the BLS uses its data on payroll workers to compute an

Income Components of Economic Output*

	Share of total 2013 output
Compensation of employees	52.4%
Corporate profits	10.1%
Rental income	3.5%
Net interest income	4.0%
Proprietor's income	7.9%
Indirect taxes less subsidies	6.5%
Depreciation	15.6%
Total	100.0%

Source: Bureau of Economic Analysis.

* Also often referred to by the BEA and others as gross value added.

Other Measurement Challenges

Measurement problems with the labor share of income do not end with proprietor's income. Fortunately, none substantially alter the conclusions reached regarding the postwar evolution of the labor share.

Government and housing sectors. In the national accounts, the value added by the government is almost exclusively assigned to labor. The only portion to accrue to capital is the consumption of fixed capital, or depreciation. Yet, there is no question that capital is a factor of production for government services, so the BEA is clearly overstating the labor share in this sector. The national accounts also arbitrarily assign to labor all the output from owner-occupied housing — except that in this case it is assigned to capital income.* There is again no question that housekeeping demands quite a bit of labor, so we are now understating the labor share in this sector. The BLS, as well as most researchers, computes the labor share of income excluding both sectors.

Indirect taxes. Most U.S. states and even some cities levy sales taxes. Thus, the government captures part of the income generated by output without being a factor of production. It is fortunately quite straightforward to compute the fraction of the sales tax “paid” by workers and investors: We just need to impute the observed split between labor and capital income for the after-tax income.

Depreciation. The depreciation, or consumption of fixed capital, is merely compensation for the physical wear and tear of capital. It is thus naturally assigned exclusively to capital income. However, it is fair to point out that workers are also subject to wear and tear! Yet, there is no entry for labor depreciation in the national accounts.

* Without a market for government goods and services, the BEA must rely on cost data to value government output.

average hourly wage. The BLS then assumes that a self-employed worker would pay himself or herself the implicit wage rate. Then, using data on hours worked by self-employed workers, it obtains a measure of the labor compensation for self-employed individuals simply by multiplying the average hourly wage by the number of hours worked by the self-employed. The result is then assigned to labor income. The rest of the proprietor's income is considered capital income.⁵

Figure 1 plots the BLS's *headline* labor share at an annual frequency from 1950 to 2013.⁶ Up until 2001, the labor share displayed some ups and downs, and perhaps a slight downward trend, but it never strayed far from 62 percent. From 2001 onward, though, the labor share has been steadily decreasing, dropping below 60 percent for the first time in 2004 and continuing its fall to 56 percent as of 2014.⁷

An alternative measure. Michael Elsby, Bart Hobijn, and Aysegül Sahin have pointed out that some of the fall in the labor share in the past 15 years is due to how the

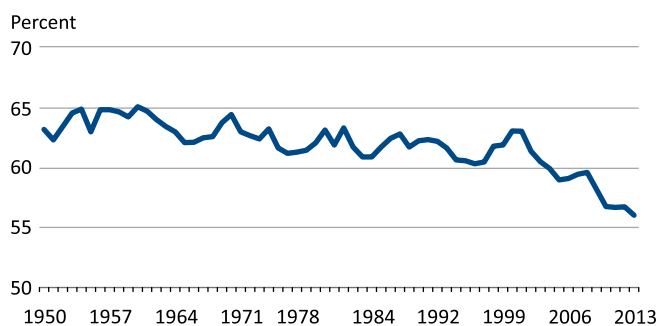
BLS splits proprietor's income. Indeed, until 2001, the BLS's methodology assigned most of proprietor's income to the labor share, a bit more than four-fifths of it. Since then, less than half of proprietor's income has been classified as labor income.

How important is this shift? It is fortunately very easy to produce an alternative measure of the labor share in which a constant fraction of proprietor's income accrues to labor. Setting that fraction to its historical average prior to 2000 — 85 percent — we can figure out what would be the current labor share under this alternative assumption. Figure 2 contrasts the previous headline number against this alternative measure from 1980 onward. First, we confirm that through 2000, both the headline and the alternative measure pretty much coincide. Since 2001, though, they diverge, with the drop being noticeably smaller in the alternative measure. Indeed, this divergence suggests that at least one-third and possibly closer to half of the drop in the headline labor share is due to how the BLS treats proprietor's income.

Alternatively, we can also proceed by the centuries-tested scientific method of ignoring the problem altogether and compute the compensation or payroll share instead of the labor share. That is, we can assume that none of proprietor's income accrues to labor. This is actually a quite common approach, since detailed payroll data exist for all industries, allowing us to pinpoint which sectors of the economy are responsible for the dynamics of labor income. The compensation share is, obviously, lower than the labor share — but its evolution across time is very similar: stable until the turn of the millennium and a decline since then.

Yet another measure. There is yet another possible way to circumvent the ambiguity regarding proprietor's income. The data allow us to zoom in to the nonfinancial corporate

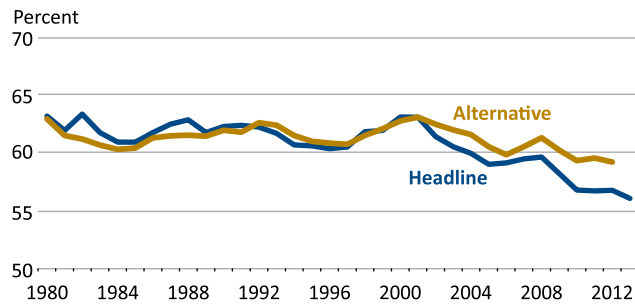
FIGURE 1
Labor's Share Began Shrinking Around 2001
Headline labor income share.



Source: Bureau of Labor Statistics.

FIGURE 2

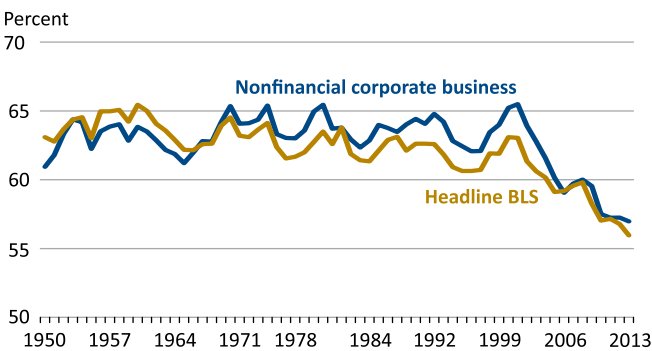
Change in Methodology in 2001 Is Evident
 Labor income share, headline versus alternative measure including all proprietor's income.



Sources: Bureau of Labor Statistics and author's calculations.

FIGURE 3

Drop Is Clear Even Without Proprietor's Income
 Labor income share, headline versus labor income share from nonfinancial corporate sector.



Sources: Bureau of Labor Statistics, Bureau of Economic Analysis, and author's calculations.

business sector. By law, corporations must declare payroll and profits separately for fiscal purposes, so there is actually no proprietor's income. The downside is, of course, that we are working with only a subset of the economy, albeit a very large one.⁸

Figure 3 plots the BLS headline measure and the labor share of income of the nonfinancial corporate sector from 1950 to the latest data available. The two series overlap for most of the period, though the headline labor share was consistently about 1 percentage point below that of nonfinancial firms from 1980 onward. In any case, the message since 2000 is unmistakable: The large drop in the headline measure is fully reflected in this alternative measure.

So, despite the inherent measurement problems, the data are clear: First, the labor share was stable from 1950

to at least near the end of the 1980s. Second, it has fallen precipitously since 2001. While the exact magnitude of the drop may be open to debate, there is no doubt that the downward trend in the labor share since 2001 is unprecedented in the data and, at the time of this writing, shows no signs of abating.

A BIT OF A MIRACLE: 1950-1987

We now take a closer look at the period in which the labor share was stable — roughly from the end of World War II to the late 1980s — by breaking it down by sector. In doing so, we will understand the logic behind the “bit of a miracle” quip. The cutoff date is necessarily 1987, since the industry classification changed in that year. Fortunately, it is also the approximate end date of the stable period for the labor share.

Since the end of WWII, the U.S. has gone through large structural changes to its sectorial composition. The most significant was the shift from manufacturing to services. In 1950, manufacturing accounted for more than two-thirds of the nonfarm business sector. By 1987, manufacturing was just half of the nonfarm business sector. Over the same period, services increased from 21 percent to 40 percent of the nonfarm business sector.⁹

The reader would not be surprised to learn that different sectors use labor and capital in different proportions. In 1950, the manufacturing sector averaged a labor share of 62 percent, with some subsectors having even higher labor shares, such as durable goods manufacturing, with a labor share of 77 percent.¹⁰ Services instead relied more on capital and thus had lower labor shares: an average of 48 percent.

Thus, from 1950 to 1987, the sector with a high labor share (manufacturing) was cut in half, while the sector with a low labor share (services) doubled. The aggregate labor share is, naturally, the weighted average across these sectors. Therefore, we would have expected the aggregate labor share to fall. But as we already know, it did not. The reason is that, coincidentally with the shift from manufacturing to services, the labor share of the service sector rose sharply, from 48 percent in 1950 to 56 percent in 1987. Education and health services went from labor shares around 50 percent to the highest values in the whole economy, close to 84 percent.¹¹ In manufacturing, the labor share was substantially more stable, increasing by less than 2 percentage points over the period.

And this is the “bit of a miracle” — that the forces affecting the labor share across and within sectors just happened to cancel each other out over a period of almost half a century.

A BIT OF A MIRACLE NO MORE: 1987-2011

I start by repeating the previous exercise, now over the period 1987 to 2011. As it had from 1950 to 1987, the manufacturing sector kept losing ground to the service sector, albeit at a slower rate. By 2011, services accounted for more than two-thirds of U.S. economic output and an even larger fraction of total employment. However, the differences in the labor share between the two sectors were much smaller by the early 1990s, and thus the shift from manufacturing to services had only small downward effects on the overall labor share.

We readily find out which part of the economy is behind the decline of the labor share once we look at the change in the labor share within manufacturing, which dropped almost 10 percentage points. Virtually all the major manufacturing subsectors saw their labor shares fall; for nondurable goods manufacturing it dropped from 62 percent to 40 percent. The labor share within the service sector kept increasing, as it had before 1987, but very modestly, only enough to cancel the downward pressure from the shift across sectors. Indeed, had the labor share of income in manufacturing stayed constant, the overall labor share would have barely budged.

Note that in one sense, the bit of a miracle actually continued from 1987 onward: As manufacturing continued to shrink, decreasing the share of income accruing to labor, services picked up the slack by increasing their share of income accruing to labor, albeit more modestly than before. What ended the “miracle” was the precipitous decline in the labor share within manufacturing.

Wages and productivity. It is worth investigating a bit further what determinants are behind the fall in the labor share within manufacturing, since it played such an important role in the decline of the overall labor share. To this end, note that the change in the labor share in a particular sector is linked to the joint evolution of wages and labor productivity. Consider a machine operator working in a factory for one hour to produce goods that will have a gross value to the factory owner of \$100. If he is paid \$60 per hour, labor’s share is approximately 60 percent. For the labor share to change, there are only two possibilities: Either the value of the goods produced must change or the hourly wage must. Conversely, for the labor share to stay constant, the value of the goods and the hourly wage have to move in unison.¹²

So which one — productivity or wages — brought down the labor share in manufacturing? Fortunately, we do

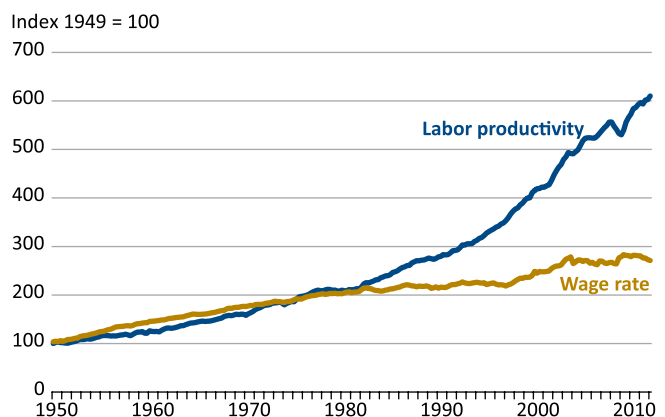
have reliable data on output, wage rates, and hours worked in manufacturing. Figure 4 displays the evolution of labor productivity (that is, output per hour) and wage rates from 1950 onward.¹³ Both series are set such that their value in 1949 equals 100.¹⁴ Once again we see two clearly separate periods. Until the early 1980s, labor productivity and wages grew at a very similar rate — if anything, the wage rate outpaced productivity, which, as described earlier, implies that the labor share in manufacturing inched up. By mid-1985, labor productivity took off, while wage growth was very sluggish. Since then, the gap between productivity and wages has kept growing, depressing the labor share.

Because an index is used to scale both series, it is a tad difficult to grasp from the figure whether labor productivity accelerated or wage rates stagnated from the 1980s onward. The answer is both things happened. In the 1980s, productivity grew at about its long-term trend rate, but wages were virtually flat, growing less than half a percentage point a year on average over the decade. Wage growth recovered in the 1990s, but productivity actually took off, further increasing the gap. Overall, though, it appears that the fall in the labor share is explained mainly by the sluggish growth of wages rather than above-trend labor productivity.

CONCURRENT PHENOMENA

What is the ultimate cause behind the decline of the labor share in the U.S.? The honest answer is that economists have several hypotheses but no definite answer yet.¹⁵ Rather than go over the sometimes-intricate theories behind these

FIGURE 4
Productivity Rose While Wage Growth Stalled
Labor productivity and wages in the manufacturing sector.



Sources: Fleck, Glaser, and Sprague (2011) and author’s calculations.

hypotheses, I will discuss the main observation or phenomenon anchoring each one.

Capital deepening. This is by far the most popular hypothesis: Workers have been replaced by equipment and software. Who has not seen footage of robots working an auto assembly line? Older readers may remember when live tellers and not ATMs dispensed cash at banks. Software is now capable of piloting planes and, even more amazingly, doing our taxes!

There is more behind this hypothesis than anecdotes. Loukas Karabarbounis and Brent Neiman document a fall in equipment prices. Lawrence Summers proposes that capital should be viewed as at least a partial substitute for labor — more and more so as technology develops. In both models, the idea is similar: Better or cheaper equipment replaces workers and redistributes income from labor to capital. The result is that production becomes more intensive in capital, which is why these theories are often referred to as *capital deepening*.

It is important to understand that the capital deepening mechanism must operate at the level of the overall economy. So, when we see a robot replace, say, five workers, we need to remember that the production of the robot itself involved workers, so we are swapping auto assemblers with robot assemblers. It is, of course, still possible that the robot tilts income toward capital, but it is not a foregone conclusion.

The main challenge to capital deepening is that if a sector is substituting robots for workers to save money or improve the quality of the good being produced, the remaining workers should therefore become more productive and, overall, the sector should be expanding. In other words, capital deepening can reduce the labor share of income, but it does so by making labor productivity accelerate rather than making wages stagnate. As we saw earlier, this does not fit the actual picture of the manufacturing sector at all.¹⁶

Income inequality. The increase in income inequality in the U.S. has lately received a lot of attention. The decline of the labor share is a force toward income inequality because capital is more concentrated across households than labor is.¹⁷

It should be noted, though, that the main driver of the increase in income inequality is not capital income but rather wages themselves, particularly at the very top of the pay ladder.¹⁸ As Elsby and his coauthors document, the increase in top wages has actually sustained the labor share. In other

words, the decline in the labor share actually understates the increase in income inequality.

An interesting question is whether whatever is driving up inequality is also driving down the labor share. Several economists have proposed that technological change is skill biased — that is, it augments productivity more for highly skilled workers than for low-skilled workers. Combined with the idea that capital helps highly skilled workers be more productive but makes unskilled workers redundant, skill bias can explain both the increase in wage inequality and the decline in the labor share.¹⁹

Let us return once more to the car manufacturer example. The robot may be replacing five unskilled workers but may require a qualified operator. The demand for unskilled workers falls, and so do their wages; but the demand for qualified operators increases, and so do their wages. So it is possible to have an increase in wage inequality while factories undergo capital deepening.

Globalization. Another popular hypothesis links the fall in the labor share with the advent of international trade liberalization. There is no question that there has been a

The decline in the labor share actually understates the increase in income inequality.

substantial increase in trade by U.S. firms in the past few decades. In particular, firms have shifted parts of their production processes to foreign countries to take advantage of cheaper inputs — which, from the perspective of a country like the U.S. that has more capital than other countries, means cheap labor. Industries that are more intensive in labor, such as manufacturing, will be more likely to outsource their production processes abroad, and thus the remaining factories are likely to be the ones that rely more on capital.

Surprisingly, there is not a lot of evidence to support this view. The main challenge to the hypothesis is that U.S. exports and imports are very similar in their factor composition. That is, were trade driving down the labor share, we would observe the U.S. importing goods that use a lot of labor and exporting goods that use a lot of capital. Instead, most international trade involves exchanging goods that are very similar, such as cars.²⁰ Another prediction of the globalization theory is that countries the U.S. exports to should see their labor shares increase and — as noted in

What About Other Countries?

To help us uncover why the U.S. labor share has evolved over time, we can look at whether economic conditions and policies in other countries had an impact on their labor shares.

In his seminal 2002 work on labor shares across countries, Douglas Gollin found enormous variation, particularly among developing economies: Ghana, for example, had a labor share below 10 percent. At the other end, Ukraine reported a labor share close to 80 percent. Moreover, there are consistent patterns with income, with poor countries being more likely than rich countries to have low labor shares.

However, Gollin pointed out some important measurement problems: Self-employment varies greatly across countries and in a systematic way with their level of development. For example, the farming and animal husbandry sectors have very low labor shares, especially in less developed countries, where self-employed workers — sometimes just a family member in charge of a small plot of land — are prevalent. At the same time, it is well documented that as a country develops, farming and self-employment both decrease. After correcting for these and other measurement issues, Gollin found that the adjusted labor shares have much lower dispersion and had no relationship with income.

Because Gollin's 2002 study used data from the United Nations National Account Statistics collected in the early 1990s, prior to the decline in the labor share in the U.S., we may ask: Have other countries experienced a fall in the labor share over the past 20 years? Loukas Karabarbounis and Brent Neiman found that, indeed, the labor share declined in most countries, with the few exceptions being some less developed economies.

This low variation across countries suggests that the decline in the labor share must be due to determinants with a global scope such as technology or trade.

the accompanying discussion, *What About Other Countries?* — it appears that the decline in the labor share is a global phenomenon.

Some studies, though, do support this hypothesis. Elsby and his coauthors find some evidence that the labor share fell more in sectors that were more exposed to imports. There is a large body of literature on the impact of trade on wage inequality that only recently has started to consider the impact on the labor share.²¹

CONCLUSIONS

Despite several measurement issues and alternative definitions associated with the labor share, the message is quite clear: The 2000s witnessed an unprecedented drop in the labor share of income. Exploring the early period, we saw that the U.S. economy had been able to accommodate the surplus workers from manufacturing only until the late 1980s. We also saw that the stagnation of wages, rather than accelerated labor productivity, has been behind the drop in the labor share from 2000 onward. The review of possible hypotheses behind the decline in the U.S. labor share was, admittedly, quite inconclusive: Economists do not yet have a full grasp of the underlying determinants. ■

NOTES

¹ A small share of income is directly captured by the government, as I discuss in more detail later.

² The Bureau of Economic Analysis offers several introductory level guides to NIPA, easily accessible at www.bea.gov.

³ Depreciation and taxes have their own set of measurement issues, but they can safely be ignored: See the accompanying discussion, *Other Measurement Challenges*, for a brief description of additional issues. See also Paul Gomme and Peter Rupert (2004) for a complete description of all labor share measurement issues.

⁴ It actually includes some other small components such as income from nonprofit institutions.

⁵ The BLS methodology is far from perfect: Some occupations may be more common among payroll workers than among self-employed workers, for example, which could lead to misleading results.

⁶ The headline labor share, also called the top-line labor share, includes all industries. See www.bls.gov/lpc/lpcmethods.htm for details on the construction of the headline measure.

⁷ The labor share also displays some weak cyclical, increasing during economic downturns. The focus of this article, though, is squarely on the long-run trends of the labor share.

⁸ In addition, we consider only nonfinancial corporations. Financial sector income is notoriously volatile and presents some measurement problems of its own — namely, that stock options and similar payments are often used as labor compensation.

⁹ Data reported as a share of total value added. The shifts in the share of employment are of similar magnitude. Over the same period, the farming sector also shrank substantially.

¹⁰ Data from Michael Elsby, Bart Hobijn, and Aysegül Sahin (2013). Labor shares are given by the share of payroll compensation over total value added.

¹¹ The increase in the labor share in services is related to William Baumol's "cost disease of services," dating back to the 1960s, that argues that productivity growth is inherently more difficult to achieve in services. See Baumol (2012) for an updated view.

¹² This calculation ignores the decomposition of proprietor's income, discussed earlier. For the manufacturing sector, though, proprietor's income is unlikely to be large.

¹³ The "wage" rate actually includes benefits and bonuses, which have become an increasing fraction of total labor compensation.

¹⁴ Following Susan Fleck, John Glaser, and Shawn Sprague (2011), Figure 4 deflates output by the implicit price index of manufacturing output and wage rates by the consumer price index.

¹⁵ Some may argue that "yet" is itself not very honest, for economists always have several hypotheses and no definitive answer for any question that is posed to them.

¹⁶ To be fair, the 1990s do fit quite well with the theory, as labor productivity did accelerate over the period. However, most of the fall in the labor share happened in the 2000s. There is also an issue of labor composition: See the discussion on page 6 of the relationship between technology and the wage gap separating skilled and unskilled workers.

¹⁷ See Margaret Jacobson and Filippo Occhino (2012) for an accessible study of the effect of the labor share on inequality.

¹⁸ See Thomas Piketty and Emmanuel Saez (2003) and Facundo Alvaredo, Anthony Atkinson, Piketty, and Saez (2013).

¹⁹ There is a long line of work on skill-biased technological change as well as capital-skill complementarity. Lawrence Katz and Kevin Murphy (1992) and Per Krusell and his coauthors (2000) are the seminal references, but both articles are quite technical. See Keith Sill's 2002 *Business Review* article for a more accessible discussion.

²⁰ The original observation was attributed to Wassily Leontief. See Daniel Treffer (1993) for an evaluation of the factor content of trade using more recent data. However, a number of studies have challenged his conclusion; see, for example, John Romalis (2004).

²¹ See Avraham Ebenstein and his coauthors (2013) for an example. For a summary of previous work, see Stephen Golub (1998).

REFERENCES

- Alvaredo, Facundo, Anthony Atkinson, Thomas Piketty, and Emmanuel Saez. World Top Incomes Database, (2013). <http://topincomes.g-mond.parisschoolofeconomics.eu/>.
- Baumol, William J. *The Cost Disease*. New Haven: Yale University Press, 2012.
- Ebenstein, Avraham, Ann Harrison, Margaret McMillan, and Shannon Phillips. "Why Are American Workers Getting Poorer? Estimating the Impact of Trade and Offshoring Using the CPS," National Bureau of Economic Research Working Paper 15107 (2013).
- Elsby, Michael W.L., Bart Hobijn, and Aysegul Sahin. "The Decline of the U.S. Labor Share," Brookings Panel on Economic Activity (September 2013).
- Fleck, Susan, John Glaser, and Shawn Sprague. "The Compensation-Productivity Gap: A Visual Essay," *Monthly Labor Review* (January 2011).
- Gollin, Douglas. "Getting Income Shares Right," *Journal of Political Economy*, 110:2 (2002), pp. 458–474.
- Golub, Stephen. "Does Trade with Low-Wage Countries Hurt American Workers?" Federal Reserve Bank of Philadelphia *Business Review* (March/April 1998).
- Gomme, Paul, and Peter Rupert. "Measuring Labor's Share of Income," Federal Reserve Bank of Cleveland Policy Discussion Paper 7 (November 2004).
- Jacobson, Margaret, and Filippo Occhino. "Labor's Declining Share of Income and Rising Inequality," Federal Reserve Bank of Cleveland *Economic Commentary* 2012–13 (2012).
- Karabarbounis, Loukas, and Brent Neiman. "The Global Decline of the Labor Share," *Quarterly Journal of Economics* (2014), pp. 61–103.
- Katz, Lawrence, and Kevin Murphy. "Changes in Relative Wages, 1963-1987: Supply and Demand Factors," *Quarterly Journal of Economics*, 107:1 (1992), pp. 35–78.
- Keynes, John Maynard. "Relative Movements of Real Wages and Output," *Economic Journal*, 49 (1939), pp. 34–51.
- Krusell, Per, Lee Ohanian, Jose-Victor Rios-Rull, and Giovanni Violante. "Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis," *Econometrica*, 68:5 (2000), pp. 1,029–1,053.
- Leontief, Wassily W. "Factor Proportions and the Structure of American Trade: Further Theoretical and Empirical Analysis," *Review of Economics and Statistics*, 38 (1956), pp. 386–407.
- Piketty, Thomas, and Emmanuel Saez. "Income Inequality in the United States 1913–1998," *Quarterly Journal of Economics* (2003), pp. 1–39.
- Romalis, John. "Factor Proportions and the Structure of Commodity Trade," *American Economic Review*, 94:1 (2004), pp. 67–97.
- Sill, Keith. "Widening the Wage Gap: The Skill Premium and Technology," Federal Reserve Bank of Philadelphia *Business Review* (Fourth Quarter 2002).
- Summers, Lawrence. "Economic Possibilities for Our Children," *NBER Reporter*, 4 (2013), pp. 1–6.
- Trefler, Daniel. "International Factor Price Differences: Leontief Was Right!" *Journal of Political Economy* 101:6 (1993), pp. 961–987.