

MICHAEL W. L. ELSBY

University of Edinburgh

BART HOBIJN

*Federal Reserve Bank of San Francisco,
VU University Amsterdam, and Tinbergen Institute*

AYŞEGÜL ŞAHİN

*Federal Reserve Bank of New York**

The Decline of the U.S. Labor Share

Prepared for Brookings Panel on Economic Activity, September 19-20, 2013.

This version: October 18, 2013.

ABSTRACT Over the past quarter century, labor's share of income in the United States has trended downwards, reaching its lowest level in the postwar period after the Great Recession. Detailed examination of the magnitude, determinants and implications of this decline delivers five conclusions. First, around one third of the decline in the published labor share appears to be an artifact of statistical procedures used to impute the labor income of the self-employed that underlies the headline measure. Second, movements in labor's share are not a feature solely of recent U.S. history: The relative stability of the aggregate labor share prior to the 1980s in fact veiled substantial, though offsetting, movements in labor shares within industries. By contrast, the recent decline has been dominated by trade and manufacturing sectors. Third, U.S. data provide limited support for neoclassical explanations based on the substitution of capital for (unskilled) labor to exploit technical change embodied in new capital goods. Fourth, *prima facie* evidence for institutional explanations based on the decline in unionization is inconclusive. Finally, our analysis identifies offshoring of the labor-intensive component of the U.S. supply chain as a leading potential explanation of the decline in the U.S. labor share over the past 25 years.

* The views expressed in this paper solely reflect those of the authors and not necessarily those of the Federal Reserve Bank of New York, Federal Reserve Bank of San Francisco, nor those of the Federal Reserve System as a whole. This version is based on data through July 28, 2013. We would like to thank Mary Amiti, Susanto Basu, Mary Daly, John Fernald, Chad Jones, Andy Levin, Sylvain Leduc, Brent Neiman, Vasia Panousi, Joe Tracy, Robert Zymek, and the editors for their many suggestions and comments. We are grateful to Yifan Cao for his excellent research assistance.

Ever since Kaldor (1957, 1961) documented his growth facts, the constancy of the share of income that flows to labor has been taken to be one of the quintessential stylized facts of macroeconomics.¹ After several decades of modest variation, prominent measures of labor's share in the United States have declined significantly. The headline measure published by the Bureau of Labor Statistics (BLS) historically fluctuated around a mean of close to 64 percent from the immediate postwar period to the mid-1980s. Thereafter, aside from a brief surge surrounding the tech bubble at the turn of the 21st century, this measure has displayed a downward trend, averaging around 58 percent in recent years, 6 percentage points below the level that prevailed during the first four decades of the postwar period.

In this paper, we examine the magnitude, determinants, and implications of the movements in the U.S. labor share over the past 25 years. Our paper follows in the footsteps of an extensive literature that has investigated fluctuations and trends in labor's share dating back to the first half of the 20th century.² We address three main themes. First, we identify the sources of income, and the underlying industry-level trends that account for decline of the U.S. labor share. Second, we consider possible explanations for the decline. Finally, we reflect on whether the recent decline warrants a major rethink of the way the labor share is used by macroeconomists.

Section I documents the measurement of the headline labor share published by the BLS, and the role played by each of its constituent income sources in its decline. We show that most of the recent downward trend in the labor share has its origins in reductions in the compensation of payroll employees as a fraction of gross value added, what we shall refer to as the "payroll share." However, the decline in the share of the remaining source of labor income, that of the self-employed, is shown to be overstated in the headline measure. This is constructed under the assumption that average wages among the self-employed are the same as those of payroll employees. We provide evidence that suggests that this assumption induced the headline measure to imply a negative capital share among the self-employed during the 1980s, thereby overestimating labor's share and casting doubt on subsequent trends in that series. Two alternative measures proposed in the early work of Kravis (1959) have less extreme implications for the returns to capital

¹ It is important to realize, however, that the proposed stability of the labor share is as stylized now as it was controversial when Kaldor classified it as a prototypical growth fact. Kaldor's claim was met by an extensive literature on measuring the movements in labor's share during the first half of the 20th century.

² For a comprehensive list of this early literature, see Gallaway (1964), footnote 1.

among the self-employed. Comparison of these two alternative measures with the headline series provides the basis of our conclusion that around a third of the decline in the headline measure is a symptom of the method used to impute self-employment income.

Once account is taken of the issues surrounding the measurement of the labor income of the self-employed, almost all of the remaining decline in the labor share is accounted for by the decline in the payroll share. For this reason, and because the payroll share is comparatively straightforward to measure, and can easily be disaggregated, we focus on understanding the movements of the payroll share in the remainder of the paper.

In Section II we address the sectoral origins of the decline in the payroll share using disaggregated industry-level data. We find that, viewed from a sectoral perspective, movements in labor's share are not a feature solely of recent U.S. history: The relative stability of the aggregate labor share prior to the 1980s in fact veiled substantial, though offsetting, movements in labor shares within industries. The shift from manufacturing to services over this period served to depress the aggregate payroll share, as labor shares in manufacturing exceeded those in services. This effect was undone, however, by substantial rises in payroll shares within industries, especially healthcare, prior to the 1980s.

By contrast, these coincidental offsetting effects have not been mirrored since the late 1980s, the period that saw the emergence of the recent decline in the aggregate payroll share. This instead has been dominated by within-industry declines in payroll shares, particularly in manufacturing and trade as opposed to compositional shifts.

A small group of industries also accounts for the higher-frequency rise and fall in the labor share in the late 1990s and early 2000s. Consistent with the timing of the tech bubble that arose at the time, and the widespread use of stock options in employee compensation in this sector, we find that around half of the rise and fall in the aggregate payroll share between 1998 and 2003 is accounted for by the changes in payroll shares of a small set of industries that cover investment banking and the tech sectors, despite their small share in total value added.

Thus, the results in Section II reveal that aggregate movements in the payroll share owe much to underlying movements in particular industries. Understanding the evolution of the aggregate payroll share therefore requires an understanding of changes in payroll shares across sectors. This insight is confirmed by our analysis in Section III.

In that section we show that explanations for the decline in the labor share that rely solely on an aggregate perspective provide, at best, an incomplete account of the movements in the labor share. A leading candidate among these explanations is the neoclassical notion that declines in the relative price of investment goods, such as computer equipment, have induced firms to replace workers with machines. This hypothesis dates back to Hicks (1932) and highlights the role of capital-labor substitutability and capital deepening in the determination of income shares, and has recently been revived by Karabarbounis and Neiman (2013). Two corollaries of this account are not supported by U.S. data, however. First, the acceleration of the decline in the labor share in the 2000s was not accompanied by an increase in the rate of capital deepening. Second, contrary to the predictions of simple versions of the theory, growth in real wages and output per hour in fact slowed down during the 2000s rather than accelerated. We find similar tensions in other, more-elaborate neoclassical explanations, such as the role played by an increasing skill-share of the labor force in the presence of capital-skill complementarity.

In Section IV, we go beyond an aggregate production function representation and investigate the sources of the within-industry changes in payroll shares by exploiting cross-industry data. Our analysis identifies a strong relationship between increases in import exposure and declines in labor share at the industry level. Strikingly we find that this statistical relationship has the potential to account for 3.3 percentage points of the 3.9 percentage point decline in the U.S. payroll share over the past quarter century. We also revisit the capital-labor substitutability hypothesis by examining whether the industries that saw the smallest increases in equipment prices are also those that experienced the largest declines in payroll shares. Reiterating our aggregate analysis, we find little support in the cross-industry data for this explanation. *Prima facie* evidence for another potential explanation—the decline in the bargaining power of workers due to deunionization—is inconclusive: there is a statistically imprecise relation between cross-industry changes in unionization rates and sectoral declines in payroll shares that includes both potentially large (and small) effects within conventional confidence intervals.³

³ A small but growing literature has focused on understanding trends in the labor share across countries. In an early example, Blanchard (1997) focuses on the declines in labor shares in continental Europe during the period 1980 to 1995. Blanchard identifies two potential explanations: the first is a shift in the distribution of rents from workers to firms; the second is the adoption of technologies that use less labor and more capital. His empirical results provide weak support in favor of the second explanation. More recent studies, such as Bentolila and Saint-Paul (2003), Guscina (2006), Harrison (2005), and Jaumotte and Tytell (2007), analyze the role of several factors in accounting for cross-country movements in labor shares. While there is no consensus on the drivers of these changes, technological progress, globalization, and openness to trade emerge as potential leading explanations.

In Section V we conclude by considering how these facts and explanations affect the way macroeconomists should think about the labor share. We first show that, in spite of the recent movements in the labor share, the assumption of (approximate) constancy of the labor share is still useful in many applications. For example, estimates of TFP growth are almost unaffected by the recent variation in factor shares observed in the United States. Second, we argue that theories and applications that use the labor share as an indicator of the state of the business cycle are affected by the decline. Such theories implicitly assume the labor share fluctuates around a constant mean; an assumption rendered invalid by the decline. Moreover, many components of the compensation captured in the labor share, like self-employment income and stock options, are not part of firms' marginal costs. This makes the labor share a highly imperfect measure of cyclical marginal cost pressures that drive inflation. Finally, we illustrate why one has to be careful to interpret the decline in the share as showing increased inequality between capital(ists) and workers. We show that the labor share in the United States has in fact been *buoyed up* increasingly by the rising income shares of very rich workers, like CEOs for example. Thus, the decline in the labor share conceals, rather than exposes, most of the large increases in inequality that have emerged in recent decades.

Of course, our conclusions are based on the recent movements of the labor share relative to its history. If the recent steep decline continues then this would lead to a more drastic reconsideration of the use of the labor share in empirical macroeconomics than we advocate. With this in mind, we conclude this paper by briefly discussing what forces might drive movements in labor's share in the foreseeable future.

I. Measurement and the impact of self-employment income

A first priority of any analysis of the evolution of the labor share is to document its empirical behavior. We begin by reviewing the path of the most commonly cited headline measure of the labor share, namely the labor share for the U.S. nonfarm business sector, which is published as part of the BLS Labor Productivity and Costs release.⁴ In this section, we contrast the behavior of this

⁴ The results in this paper are based on data through July 28, 2013 and do not include the 2013 benchmark revisions of the National Income and Product Accounts, which affect the level and path of the labor share in two ways. First, they add part of intangible expenditures to value added, due to the capitalization of intangible investments. This increase in value added shows up as an increase in net operating surplus and, by construction, a reduction in the labor share. This increase is partly offset by an increase in measured compensation due to a change in the way pension benefits are accounted for. In the end, these revisions had little effect on the trends in the labor share.

published series with several alternative measures of the labor share. Our findings indicate that the headline measure has overstated the decline in the labor share because of the way self-employment income is imputed in the construction of the headline number.

I.A. Evolution of the headline labor share measure

The evolution of the headline labor share measure during the postwar period is depicted by the black line in Figure 1, and some related summary statistics are provided in Table 1. Four observations are apparent from these. First, the labor share exhibited no obvious trend during the first four decades of the postwar era. From 1948 through 1987 it hovered around a mean of 63.6 percent. Second, one can discern a trend decline in the share since the late 1980s. Third, the share spiked in the late 1990s and early 2000s. Finally, the trend decline that started in the late 1980s has accelerated over the last decade with the labor share receding from its high in 2001 to a historic postwar low in the wake of the Great Recession. It averaged 58.3 percent over the period 2010 to 2012. The difference between the 1948-1987 and 2010-2012 period averages implies that the share of income that flows to workers has declined by 6.4 percentage points.

The methodology used to construct this headline measure is but one of a number of alternative approaches to measuring labor's share, however. All these measures have one thing in common: They capture the ratio of a measure of aggregate labor compensation to an analogous measure of all factor income generated, which equals gross value added. Thus, the labor share, which we denote by λ , is equal to the ratio of labor compensation to gross value added

$$\lambda \equiv \frac{WL}{PY}, \quad (1)$$

where W is the average compensation per hour worked, L is the total number of hours worked, P is the price and Y is the quantity produced. While this formulation is conceptually simple, neither the numerator nor the denominator is straightforward to measure. The measurement of labor compensation, WL , turns out to be particularly important for the decline in λ .⁵

⁵ As far as the denominator is concerned, the measurement of gross value added involves the choice of what are considered intermediate inputs and what are capital goods. A change in the classification of some of these expenditures, as in the 2013 benchmark revisions of the NIPA, results in a change in the level of gross value added. It also involves the choice which sectors to include and which not, and whether the denominator is constructed using data from the expenditure or income sides of the national accounts (Grant 1963). For the postwar period the latter turns out to be immaterial because the statistical discrepancy was relatively small during that period.

Labor compensation is the sum of payments to two groups of workers. The first group is comprised of workers who are on the payrolls of their employer. Their compensation consists of their wages and salaries, as well as supplements to them in the form of employer contributions to pension and insurance funds and to social insurance. We denote the compensation of these payroll employees by $W^P L^P$. Measurement of $W^P L^P$ is relatively straightforward, being based primarily on employer payroll records from the Quarterly Census of Employment and Wages, which covers 98 percent of payroll jobs in the United States.⁶ For this reason, it is often referred to as the “unambiguous” part of labor compensation (Gomme and Rupert 2004). In what follows, we will refer to the share of value added paid to these employees as the *payroll share*, and denote it by $\lambda^P = W^P L^P / PY$.⁷ The dark-shaded area in Figure 1 reflects this payroll share.

The second group of workers consists of those who work for themselves rather than for an employer. The income of these self-employed individuals⁸ reflects both returns to their work effort as well as returns to the business property they invested in. As Krueger (1999) points out, isolating the component of self-employment income that accrues to labor (as opposed to profits or other factors of production) is fundamentally ambiguous. Consequently, there is no direct measure of (hourly) *labor* compensation for the self-employed. Thus, measures of labor’s share must take a stand on this to calculate the second component of labor compensation.⁹ In addition, as Hurst, Li, and Pugsley (2012) argue there is substantial underreporting of self-employment income in household surveys.¹⁰

We denote the imputed component of labor compensation paid to the self-employed by $W^S L^S$. The light-shaded area in Figure 1 is the fraction of value added that is attributed to self-employment income in the headline measure of the labor share. Self-employment income in this headline

⁶ Estimates for the most recent quarter are based on the more timely survey data from the Current Employment Statistics program of the BLS.

⁷ It is also referred to as the “employee compensation share” (Gollin 2002).

⁸ Also known as proprietors’ income, or non-corporate business income.

⁹ In addition, from a practical perspective, there are known measurement problems associated with reported self-employment income. Chief among these is that available measures are thought to be subject to substantial underreporting. Based on Internal Revenue Service estimates of “tax gaps” for the 2001 tax year, Slemrod (2007, Table 1) reports that as much as 57 percent of nonfarm proprietors’ income is not reported on tax returns when it should have been.

¹⁰ Changes in underreporting of self-employment income can potentially affect the evolution of the labor share. However, Figure 4 in Hurst, Li and Pugsley (2010) suggests that the timing of changes in underreporting does not coincide with the evolution of the labor share. According to their calculations, underreporting was high during the early 1980s, fell between the late 1980s and mid 1990s, increased once again during the late 1990s and early 2000, but still remains below its level in the early 1980s when the labor share was higher.

measure is imputed under the assumption that average hourly compensation for the self-employed is the same as for those on payroll. That is, $W^S = W^P$ for all quarters in the postwar period. Under this assumption, the labor share can be constructed by rescaling the payroll share by one plus the ratio of the hours of the self-employed and payroll employed,

$$\lambda \equiv \left(1 + \frac{L^S}{L^P}\right) \lambda^P. \quad (2)$$

The requisite underlying source data on hours worked for replicating the BLS methodology are available back to 1964. The dashed line in Figure 1 demonstrates that, using these data, we are able to replicate closely the published labor share under this assumption. The share of hours of the self-employed in total hours, L^S/L , has declined steadily from about 14 percent in 1948 to 8.5 percent in 2012. This is why Figure 1 indicates a declining wedge between the labor and payroll shares during the postwar period. In spite of the relatively small share of self-employment hours, the treatment of self-employment income plays an important role in the recent behavior of the evolution of the labor share plotted in Figure 1.

I.B. The role of the treatment of self-employment income

The assumption of equal hourly compensation for payroll employees and the self-employed used to construct the headline labor share measure is only one way to impute the labor income of the self-employed. In this subsection we examine the behavior of a set of alternative measures of the labor share based on different assumptions on the measurement of $W^S L^S$. This approach was pioneered by Kravis (1959, Table 1).

The basic problem that each of these measures addresses is that the different sources of factor income reported in the national accounts do not map directly into parts attributable to capital and labor. Figure 2 and Table 1 summarize the distribution of all factor income generated by the nonfarm business sector across five broad categories. The first is payroll compensation and unambiguously reflects payments to labor. The second category is proprietors' income, which as discussed above reflects a mixture of the labor and capital income of the self-employed. The remaining three categories are better classified as either direct or implicit payments to capital.¹¹

¹¹ The taxes that are part of this component are "taxes on production and imports" (TOPI). They include Federal excise taxes and customs duties, state and local sales taxes, property taxes (including residential real estate taxes), motor vehicle licenses, severance

Measures of the labor share thus have to take a stance on the part of proprietors' income that reflects payments to labor.

Kravis (1959) introduced four different assumptions that can be used to cut the Gordian knot of estimating the labor income of the self-employed, and we provide an update of his results. The first assumption corresponds to that underlying the headline measure—that is, equal wages for self-employed and payroll employed. This is what Kravis calls the “Labor basis” measure of self-employment income. For reference it is plotted as the “labor share” in Figure 2.

Kravis' second measure attributes all of proprietors' income (exclusive of capital consumption allowances, CCA, and inventory valuation adjustments, IVA) to the labor input of the self-employed. We call this the “all-to-labor basis” measure. The labor share that one obtains using this measure is represented by the two lightest shaded areas in Figure 2. While this clearly is an overstatement of the labor compensation of the self-employed—it assumes after all that none of proprietors' income flows to capital—it is nonetheless instructive.

A comparison of this “all-to-labor basis” share with the published labor share is particularly revealing. Figure 2 shows that the share of proprietors' income included in the labor share implicitly increased from around 80 percent in 1948 to 100 percent in 1981. From 1981 through 1991, the level of imputed self-employment income in the headline labor share measure even *exceeded* the level of proprietors' income. Crucially for recent trends in the published share, between 1987 and 2012 the share of proprietors' income attributed to labor implicit in the headline measure has decreased from 121 percent to a postwar low of 55 percent. Consequently, as can be seen by combining the first two rows of Table 1, at 4.5 percentage points the decline in the “all-to-labor basis” is substantially smaller than the 6.4 percentage point decline in the headline measure. Moreover, rows 1 through 5 of the table reveal that only a small part of the decline in the “all-to-labor basis” share is due to an increase in the share of corporate profits. Instead, the bulk of the decline can be traced back to an increase in the share of capital services as reflected in rental and interest income, as well as capital depreciation.

The fact that the headline labor share measure exceeded the “all-to-labor basis” measure from 1981 to 1991 has the pathological implication that the capital share, and hence the marginal product

taxes, and special assessments. Some studies, like Gomme and Rupert (2004), classify taxes as ambiguous income that reflects a mixture of payments to capital and labor as well. The replication files provided with this paper also include a measure of the labor share that divides these taxes up into payments to capital and labor. This does not matter much for the implied decline in the share.

of capital, were in fact significantly *negative* during the 1980s in the proprietors' sector. By the same token, the large movements since the 1980s in the capital share of proprietors' income implicit in the published series imply very large swings in return to capital among proprietors. These observations indicate that the headline number significantly overstated the actual labor share during the 1980s, and calls into question subsequent trends in that series. To get a sense of this, it is useful to compare the headline number with the remaining two measures proposed by Kravis (1959).

The first of these is based on the assumption that the returns to capital, as captured by its user cost, are the same for the capital used by the payroll employed and the self-employed. Kravis refers to this as the "Asset basis" measure of self-employment income. Conceptually, this approach is very similar to the way self-employment labor income is inferred in the BLS Multifactor Productivity (MFP) data, and so we present that estimate here.¹² This measure has, for example, been used by Fernald (2012) to construct his real-time growth accounting dataset. The second of the remaining Kravis measures instead assumes that the labor share in entrepreneurial income is the same as that for the overall economy, what Kravis names the "Economy-wide basis" measure. Gomme and Rupert (2004) favor this measure, for example.

Figure 3 compares these two alternative measures with the published headline measure, and Table 1 includes some related summary statistics. Three observations stand out. First, from 1948 through 1987 the published measure is, on average, more than 1.5 percentage points higher than the two alternative measures. Second, aside from a period during the 1980s, the two alternative measures lie very close to each other. Third, since the mid-1990s all three measures have converged. Rows 6 through 8 of Table 1 summarize the impact of these three observations on the measured declines of the labor share. The decline in each of the alternative labor share measures is at least 1.7 percentage points smaller than the 6.4 percent decline in the headline measure.

The upshot of these comparisons is that around one third of the decline in the headline measure of labor's share appears to be a byproduct of the methods employed by the BLS to impute the labor income of the self-employed. Alternative measures that have less extreme implications regarding the return to capital among proprietors are more consistent with one another, and indicate a more modest decline.

¹² Since Kravis' (1959) analysis predates Jorgenson's (1963) derivation of the user-cost equation, Kravis' "Asset basis" measure is different than the one from the MFP release, which is user-cost based.

I.C. Reevaluating the headline measure

Since the published share is constructed under the assumption of identical payroll and self-employment wages, and taking the alternative series in Figure 3 at face value, the recent convergence of these measures is suggestive of a sharp rise in the average hourly compensation of proprietors relative to the payroll employed since the late 1980s. Such large trend shifts in relative wages of the self-employed therefore potentially play a crucial role in the decline of the headline labor share measure. A key question, then, is what might explain these trend shifts.

The most compelling evidence emerges from comparisons of the evolution of the income distributions within, respectively, the payroll employed, and self-employed. Following Piketty and Saez (2003), we disaggregate the structure of the labor share associated with wages and salaries, and proprietors' income, by fractile of the population.¹³ This analysis can be implemented through 2010 thanks to updated data posted online by Alvaredo, Atkinson, Piketty and Saez (2013) from their World Top Incomes Database, which in turn are drawn from Internal Revenue Service tax returns statistics.

The results of this exercise are illustrated in Figure 4. The aggregate series depicted by the uppermost bold black line corresponds to the sum of the compensation and proprietors' income areas in Figure 2, the "all-to-labor basis" measure of the labor share. The lower dashed black line corresponds to payroll compensation as a fraction of value added, λ^P in the previous notation. The shaded regions that comprise these two aggregates depict the portion of the labor share from each source that is accounted for by each fractile of employees and proprietors respectively.

Figure 4 demonstrates that the share of income accounted for by both payroll wages and salaries and by proprietors' income has been buoyed up since the 1980s by substantial rises in the shares accounted for by the very top fractiles of households in the United States. Indeed, the majority of the stability of the aggregate series has been sustained by rises in the income shares of just the top 1 percent of taxpayers.

This rise in inequality is even more striking for proprietors' income than it is for payroll income. In 1948 the bottom 90 percent of employees earned 75 percent of payroll compensation. By 2010 this had declined to 54 percent. For entrepreneurial income, however, this fraction declined from 42

¹³ Strictly speaking, Piketty and Saez emphasize that the data correspond to fractiles of the distribution of "tax units," defined as a married couple, or single adult, together with their dependents, as opposed to individuals.

percent in 1948 to 14 percent in 2010. Even more starkly, over the same period the share of proprietors' income accounted for by the bottom 99 percent fell from 74 percent to 45 percent. This suggests that the sharp rise in the average hourly compensation of proprietors relative to the payroll employed since the late 1980s is related to substantial increases in income inequality among proprietors that dominate even the considerable rise in inequality witnessed among the payroll employed. Moreover, this has been driven by extreme rises in proprietors' income at the very top of the income distribution—the top 1 percent in particular.¹⁴

I.D. Summary

In this section, we have provided evidence to suggest that part of the trend decline observed in the published headline measure of the labor share in the United States is spurious. We have highlighted the important role played by the methods used to impute the labor income of the self-employed, in particular the assumption of equal hourly compensation between employees and the self-employed. The two most commonly used alternative measures of the labor share—the “Economy-wide basis” and “Asset basis” measures—exhibit more modest trend declines, and provide similar depictions of movements in labor's share. Moreover, rises in compensation at the very top of the distribution of proprietors have been even more extreme than among employees, suggesting that the average hourly compensation of the self-employed has soared in recent decades relative to the payroll employed, violating a key assumption underlying the headline measure.¹⁵

It is important to note, however, that the measures of the labor share we consider also have much in common. In terms of long-run movements, Table 1 shows that all measures indicate a substantial decline in labor's share since the late 1980s. In terms of short-term fluctuations, Figures

¹⁴ Of course, this evidence is merely suggestive. For one, we do not have measures for hours worked for the IRS tax units. However, the trends in inequality are of a magnitude too large to be offset by similar differential trends in hours worked. Moreover, this increase in inequality among entrepreneurial income also could reflect an increase in inequality in income flows from capital held by entrepreneurs rather than in labor. Inequality, especially for proprietors' income, might be understated in Figure 4 because, as Johns and Slemrod (2010) point out, underreporting of income is more prevalent among high income earners.

¹⁵ In light of this, we suggest that the BLS adjust its headline labor share estimate to align more closely with either of the two alternative measures considered above. A quarterly estimate of the Multifactor Productivity (MFP) share would be preferable to be consistent with the MFP data currently published by the BLS. This would imply a different path of average compensation per hour—it would grow faster after the 1980s. One could continue to produce the current compensation per hour measure and rename it “payroll compensation per hour”. In addition, the BLS could publish payroll hours as well as the payroll share in addition to the labor share. All these data already are part of current calculations but are simply not included in the official release.

1 and 3 reveal that all measures are countercyclical, with the share rising during recessions and falling during expansions.¹⁶

A convenient aspect of the common movements in the U.S. labor share measures is that they are mirrored by similar fluctuations in the payroll share (see Figure 1 and Row 1 of Table 1). Consequently, an understanding of movements in the U.S. labor share can be gleaned by accounting for the drivers of the payroll share, which is comparatively straightforward to measure, and can also be disaggregated, for example by industry. This is the task of the next section.

II. The composition of economic activity

The previous section examined the empirical behavior of the labor share from an aggregate perspective. It established that the bulk of the recent decline in the labor share can be traced back to a decline in the aggregate payroll share. This, however, veils the underlying disaggregate origins of movements in the aggregate payroll share. In this section we take a more microeconomic perspective by exploiting data on the sectoral composition of labor compensation and production. This allows us to shed light on (i) the degree of heterogeneity in the level and changes in payroll shares across sectors; (ii) systematic movements in the sectoral composition of the U.S. economy; and (iii) how the interaction of (i) and (ii) affects movements in the aggregate payroll share.

II.A. Decomposing the decline in the payroll share

Our starting point is the simple observation that the aggregate payroll share, λ^P in the above notation, is the value-added weighted sum of the payroll shares of each of the industries in the economy:

$$\lambda^P \equiv \frac{\sum_i W_i^P L_i^P}{PY} = \sum_i \omega_i \lambda_i^P, \quad (3)$$

where $\omega_i \equiv P_i Y_i / PY$ is the value-added share of sector i , and $\lambda_i^P \equiv W_i^P L_i^P / P_i Y_i$ its payroll share. It is possible to measure each of the components of equation (3) using the industry accounts from the Bureau of Economic Analysis. Data by detailed industry are available at an annual frequency from

¹⁶ A result of these long- and short-run properties is that none of the measures is estimated to be stationary around a constant mean. Dickey-Fuller tests for unit roots and KPSS tests for mean stationarity suggest that all three measures of the labor share in Figure 3 are not mean stationary.

1948 to 1987 under the SIC72 industry classification, and from 1987 onwards based on the NAICS02 coding.¹⁷

Figure 5 depicts the evolution of the cross-industry distribution of payroll shares over the postwar period. Consistent with equation (3), the distribution is weighted by value-added so that its mean replicates the path of the aggregate payroll share illustrated in Figure 1. A key impression of Figure 5 is that the magnitude of the variation in the aggregate payroll share (the mean) is dwarfed by the variation in these shares across industries. For example, in the most recent data industries at the 25th percentile have a payroll share of around 30 percent while the corresponding number for the 75th percentile is 70 percent.

In light of this cross-industry variation, it is remarkable that the aggregate share has remained as *stable* as it has given the large shift in the composition of U.S. economic activity from manufacturing to services over the postwar period. This observation echoes Keynes' (1939) sentiment that the "remarkable constancy" of the aggregate labor share is "a bit of a miracle." This miraculous constancy of the payroll share is the outcome of movements across industries offsetting one another in subtle ways. To demonstrate this, Table 2 implements a "shift-share" analysis of the change in the aggregate payroll share. Using equation (3), one can decompose changes in the aggregate share into two components:

$$\Delta\lambda^P = \underbrace{\sum_i \omega_i \Delta\lambda_i^P}_{\text{Shift}} + \underbrace{\sum_i \Delta\omega_i \lambda_i^P}_{\text{Share}}. \quad (4)$$

The "shift" component summarizes the contributions of changes in within-sector payroll shares $\Delta\lambda_i^P$ to the change in the aggregate payroll share. The "share" component reports the contributions of changes in sectoral composition $\Delta\omega_i$. Due to a change in industry coding in 1987, we perform the analysis separately before and after this break. For the sake of brevity, we present the results obtained at the major sector level.

The first three columns of Table 2 confirm the familiar observation that the composition of U.S. economic activity has shifted away from manufacturing towards services. The value-added share of

¹⁷ We use detailed data for 59 SIC72 industries and 60 NAICS02 industries. The GDP by industry data do not distinguish between nonprofit institutions serving households (NISHs) and businesses. As a result the data include the compensation of NISH employees. To make the data for our analysis as comparable to the nonfarm business sector as possible, we took compensation for all industries excluding farms from the GDP-by-industry data and then rescaled both compensation and value added for each industry such that they aggregate to that of the nonfarm business sector. Details are available in the online appendix.

“Manufacturing and trade, transportation and utilities” declined by a total of 18.1 and 10.5 percentage points respectively between 1948 and 1987, and between 1987 and 2011. This loss is almost exactly mirrored by the gains of Financial activities, Professional and business services, and Education and health services. The impact of the changing composition of economic activity on the payroll share is summarized in the share-part of the shift-share analysis.

Panel A of Table 2 reports results for the earlier period, 1948 to 1987. During that period, payroll shares in the services sector tended to be smaller than in manufacturing. Therefore, everything else equal, the shift in activity from manufacturing to services would have resulted in a 2.7 percentage point decline in the labor share over that period, the “share” component of Table 2A. This decline did not manifest itself, however—the payroll share in fact increased by 0.9 percentage points between 1948 and 1987. The reason is that the effect of the change in the composition of economic activity was offset by an increase in within-sector payroll shares, particularly in the service sectors, the “shift” part of Table 2A. Absent coincidental movements in sectoral composition, the aggregate payroll share would have risen by 3.5 percentage points between 1948 and 1987. It is these coincidental movements that resulted in the “miraculous” relative constancy of the labor share in the first four decades of the postwar period.

Such coincidental movements did not take place between 1987 and 2011. This can be seen from Panel B of Table 2. This demonstrates that, between 1987 and 2011, the transition from manufacturing to services has had little impact in aggregate.¹⁸ The reason is that payroll shares by major sector have converged in recent decades, so that compositional shifts have been broadly offsetting.¹⁹ What has driven the decline in the aggregate payroll share has been the decline in the fraction of income paid to workers in manufacturing, trade, transportation, and utilities. These are the sectors with a large negative contribution to the “shift” part of the shift-share analysis.

The results in Table 2 provide an important insight, namely that the recent trend decline in the aggregate payroll share is due mostly to declines in payroll shares within particular industries. Thus, understanding the sources of the decline in the aggregate payroll share requires understanding the declines in the payroll shares of individual industries.

¹⁸ A point also made by Karabarbounis and Neiman (2013). Estrada and Valdeolivas (2012) show the same is true for France, Germany, Italy, Japan, Spain, and the U.K.

¹⁹ What has increased is the within-major-sector variation in labor shares across detailed industries. This is why Figure 5 does not show a large decline in the cross-industry variation in payroll shares.

II.B. Cyclical fluctuations in the payroll share

Although our primary focus is on trend movements in the labor share, our analysis of disaggregated industry data also uncovers striking regularities in movements of the payroll share at higher frequencies.

The first regularity is that the *cyclical* movements in the aggregate payroll share apparent in Figure 1 are accounted for entirely by cyclical fluctuations within industries; changes in industry composition are immaterial for these high-frequency fluctuations. To establish this, Table 3 uses equation (4) to motivate a decomposition of the *variance* of changes in the aggregate payroll share (as opposed to the decomposition of long differences in Table 2). Since the variance is dominated by high-frequency variation, this provides a sense of the sectoral origins of the cyclical variation in the payroll share. Table 3 reveals that, for both subsamples, the “shift” part of the variance decomposition contributes essentially for all of the variance. Additionally, for all sectors, except mining (which contains oil and gas extraction), the contribution of the shift component is positive. This indicates that for all these sectors changes in individual payroll share move in the same direction as in the aggregate.

As noted in the previous section, a conspicuous aberration from the recent trend decline in λ^P is the large rise and fall observed at the turn of the 21st century. Given its timing, it is natural to conjecture that this episode has its roots in the tech bubble that emerged at the time. In particular, the analyses of Mehran and Tracy (2001) and Himmelberg, Mahoney, Bang, and Chernoff (2004) suggest that the extensive use of stock options in the compensation of employees in the tech and investment banking industries substantially affected compensation and profits measures in the national accounts between 1998 and 2003. As Moylan (2008) explains, stock options generally are counted as part of compensation only when employees exercise them.²⁰ Many tech and investment banking employees decided to exercise their options at the height of the tech bubble, propping up the payroll shares in tech and investment banking.²¹

To discern the effect of these stock options, we calculate the contributions to the aggregate payroll share accounted for by changes in payroll shares of a small set of industries that cover

²⁰ Unfortunately, separate data on compensation in the form of stock options is not available in the NIPA, because the source data from the IRS does not distinguish that part of compensation from regular wages and salaries.

²¹ For example, the payroll share in Computer and electronic products (334) increased from 75 percent in 1996 to 112 percent in 2000, and the payroll share in Information and data processing services (514) peaked at 181 percent in 2000.

investment banking and the tech sectors.²² Together, these industries account for about 7 percent of nonfarm business value added. The results are remarkable. Figure 6 plots the annual percentage point change in the aggregate payroll share from 1988 onwards, together with the contribution of investment banking and tech. Despite their small share in total value added, the contributions of these sectors are substantial. Around half of rise and subsequent fall in the aggregate payroll share between 1998 and 2003 can be attributed to changes in the payroll shares in these sectors alone. A similar story holds for the period surrounding the Great Recession, though with a greater role played by investment banking than by tech sectors, as one would expect.

Thus, mirroring the origins of the trend decline in the payroll share, movements in the share in individual industries also drive a large part of cyclical movements in the aggregate. The striking lesson of Figure 6 is that even very tightly defined sectors can experience such large swings in payroll shares that they leave a clear imprint on the evolution of these shares in the aggregate.

III. The structure of aggregate technology

So far, we have limited ourselves to a quantitative description of what accounts for the movements of the U.S. labor share in the nonfarm business sector. In this section and the next we consider possible explanations for declines in the labor share. Our focus in this section is on models of aggregate technology that potentially give rise to such movements.

To fix ideas, we begin by investigating the simplest possible model in which movements in the labor share can be understood. Our point of departure, then, is the neoclassical theory of distribution for an economy comprised of firms that operate in competitive product and factor markets and are endowed with a given technology. The historical stability of labor's share has informed the widespread assumption in much of macroeconomics of a Cobb-Douglas technology. Thus, the focus of this section will be to assess the promise of alternative structures of technology in their ability to account for recent movements in labor's share.

Before we delve further into the analysis, it is worthwhile to note the economic significance of such a neoclassical explanation, namely that it suggests that movements in factor shares are

²² The detailed industries for the Tech sector include Computer and electronic products (334); Electrical equipment, appliances, and components (335); Publishing industries (includes software) (511); Information and data processing services (514); Computer systems design and related services (5415). For investment banking, we focus on Securities and investments (523).

potentially informative on the structure of production and the nature of technical progress in the economy, two fundamental ingredients of macroeconomics.

III.A. The role of capital-labor substitutability

We start by investigating simple aggregative models. Imagine an economy endowed with a constant-returns-to-scale aggregate production function $Y = F(A_K K, A_L L)$. Here, K is capital, L is labor (hours worked), and technical progress may be either capital-augmenting (A_K) or labor-augmenting (A_L). Note that this nests both Harrod-neutral (purely labor-augmenting, $A_K \equiv 1$) and Hicks-neutral (equally capital- and labor-augmenting, $A_K \equiv A_L$) technical change as special cases.

These assumptions on the structure of production allow one to express labor's share as

$$\lambda \equiv \frac{wL}{Y} = \frac{w}{A_L f(k)} \quad (5)$$

where w is the real wage, $k \equiv (A_K K)/(A_L L)$ is the ratio of efficiency units of capital to efficiency units of labor, and $f(k) \equiv F(k, 1)$. Furthermore, if real wages w and the rental rate on capital r are determined competitively, then

$$w = \frac{\partial Y}{\partial L} = A_L [f(k) - kf'(k)], \text{ and } r = \frac{\partial Y}{\partial K} = A_K f'(k). \quad (6)$$

Thus, we can write labor's share as

$$\lambda(k) = 1 - \frac{kf'(k)}{f(k)} \equiv 1 - \alpha(k). \quad (7)$$

Due to the presence of competitive product markets, the remainder $\alpha(k)$ is of course capital's share.

Two key implications emerge from this simple model: First, labor's share is a function solely of the capital-labor ratio expressed in efficiency units, k . Second, it is determined exclusively by the structure of production embodied in $f(\cdot)$. It follows that, in this simple economy, the labor share can move over time if technology is not isoelastic (that is, not Cobb-Douglas), and the ratio of efficiency units of capital to efficiency units of labor k moves over time.

Since Hicks (1932), it has been well understood that the relationship between the labor share λ and the effective capital-labor ratio k can be linked generally to the elasticity of substitution between effective capital and labor, which we shall denote σ :

$$\sigma = \frac{d \ln k}{d \ln(F_2/F_1)} = \left[1 + \frac{1}{1 - \lambda(k)} \frac{d \ln \lambda(k)}{d \ln k} \right]^{-1} \cong 1 \text{ as } \lambda'(k) \leq 0. \quad (8)$$

Thus, the labor share $\lambda(k)$ is predicted to decrease in k whenever the elasticity of substitution σ is greater than one (and *vice versa*). Put another way, whenever effective labor and capital are more substitutable than implied by a Cobb-Douglas production function, the labor share will decline in the effective capital-labor ratio k .

Eighty years after Hicks' insight, Karabarbounis and Neiman (2013) revived this notion as a potential explanation of the evidence they present for secular declines in the global labor share. Based on cross-country variation, they infer estimates of the elasticity of substitution in the neighborhood of $\sigma = 1.25$.²³ Importantly, however, the presence of an elasticity of substitution different from one is not a sufficient condition for generating long run movements in labor's share. For example, a standard result is that purely labor-augmenting technical change will imply that the effective capital-labor ratio k will be constant along a balanced-growth path, and thus so will the labor share, regardless of the elasticity of substitution.²⁴ The key to this story is the generation of long-run movements in capital intensity; in the case of $\sigma > 1$, it is necessary for the effective capital-labor ratio k to trend upward.

Neoclassical theory suggests two potential sources for capital deepening of this sort. The first is through capital-augmenting technical progress—that is, if A_K in the above notation grows over time. A second channel is via investment-specific technical change driven by reductions in the relative price of investment associated with efficiency gains in the production of new capital goods (Greenwood, Hercowitz and Krusell 1997). Karabarbounis and Neiman (2013) highlight the latter as a particularly promising candidate explanation for the decline in labor's share.

A great virtue of the simple neoclassical framework is that it delivers a rich set of empirical predictions—on the joint evolution of labor's share, factor prices and productivity—that can be confronted with the wealth of available data to assess the ability of this explanation to account for the evolution of labor's share in the United States.

To demonstrate this, we begin with a simple observation. It is apparent from the alternative series in Figure 3 that labor's share has not trended downward throughout the postwar era. Despite

²³ By contrast, Antras (2004) and Chirinko (2008) suggest estimates of σ in the range 0.4 to 0.6.

²⁴ Movements in labor's share could arise from transitional dynamics around the balanced growth path in this case, however.

some differences across measures, the strongest evidence of a decline emerges after the 1980s. (This, of course, is why the proposed stability of the labor share has been such a robust stylized fact of macroeconomics.) Thus, for the decline of labor's share to be a manifestation of the structure of production and the nature of technical progress requires that one or both of these *changed* at some point in time. This simple observation has interesting implications for the process by which a decline in labor's share will emerge.

A first corollary is a qualitative one. As the preceding discussion demonstrates, if $\sigma > 1$, the key to a shrinking labor share must be traced to a form of capital deepening—that is, the capital-labor ratio in efficiency units k must grow. But neoclassical theory also predicts that growth in k will be associated with other adjustments in the economy. An important example in the U.S. context is that, for a given growth rate of labor-augmenting technical progress, output-per-hour $Y/L = A_L f(k)$ and real wages $w = A_L [f(k) - kf'(k)]$ will be predicted to grow faster. And, consistent with a declining labor share, the additional growth in output-per-hour will outstrip the additional growth in real wages.

Figure 7 confronts this prediction with the path of real compensation per hour and output per hour experienced in the U.S. economy. Specifically, it plots the growth rates of these variables as measured by the BLS Labor Productivity and Costs program for the nonfarm business sector.²⁵ The growth rate of the labor share, of course, is simply the difference between growth in hourly compensation and output per hour. All series are plotted as 10-year centered moving averages to smooth out high-frequency variation associated with business cycles, short-run adjustment dynamics and so forth.

The decomposition of the decline in labor's share in Figure 7 suggests that the experience of the United States in recent decades is only partly consistent with a neoclassical story based on an elasticity of substitution in excess of one together with capital deepening. Viewed through this lens, two distinct episodes are apparent. The decline in the labor share from the 1980s to the mid-1990s was characterized in part by a period of rising growth in the late 1980s during which growth in average labor productivity surpassed hourly compensation growth, consistent with the neoclassical story above.

²⁵ Bosworth and Perry (1994) point out the sensitivity of the trend properties of the real wage with average labor productivity with respect to the deflator being used for the real wage. We use the nonfarm business output deflator here. Using the CPI or PCE deflator would result in slower real wage growth since the 1980s and only strengthen the point we make.

By contrast, however, the substantial recent decline in the labor share that emerged at the turn of the 21st century appears wholly due to a slowdown in growth marked by a profound, and unprecedentedly sharp, stagnation of hourly compensation growth. This important feature of the recent U.S. experience appears hard to reconcile with a simple story based on the structure of production and technical change.

In addition to this qualitative observation, the simple theoretical analysis above also highlights quantitative implications of the timing of the decline in labor's share. Rearranging equation (8) provides a tight link between the elasticity of substitution σ and the magnitude of the response of labor's share to the effective capital-labor ratio k :

$$d \ln \lambda = -(1 - \lambda) \frac{\sigma - 1}{\sigma} d \ln k. \quad (9)$$

Evaluating the latter at a value of $\sigma = 1.25$ and a reasonable value for labor's share, say $\lambda = 0.6$, suggests a $d \ln \lambda(k)/d \ln k$ of around -0.08 . Thus, to generate an 8-percent decline in the labor share—which is about the magnitude seen in the data—it is necessary for the capital-labor ratio in efficiency units, k , to double. If, as neoclassical theory suggests, the latter is the outcome of either capital-augmenting or investment-specific technical progress, growth in the capital-labor ratio K/L will be predicted to accelerate for a given growth rate of labor-augmenting technical progress.

This suggests an additional simple check on the neoclassical account for movements in labor's share, namely to confront this prediction with the empirical behavior of modern estimates of the capital-labor ratio. Figure 8 implements this using the real-time growth accounting data recently developed by Fernald (2012).²⁶ He provides a measure of capital input constructed from an aggregate of individual measures of heterogeneous capital goods based on the perpetual inventory method. Consistent with firm optimization, the aggregate capital stock is measured as a service-flow weighted sum of the stocks of each type of capital, where the weights are based on user cost estimates for each type.

Figure 8 plots the growth rate of the capital-labor ratio implied by Fernald's measure of capital input and the BLS Labor Productivity and Costs measure of hours worked for the Business sector. Again, series are plotted as 10-year centered moving averages to highlight medium-run variation. It is clear from Figure 8 that there have been periods of elevated growth in the capital-labor ratio.

²⁶ The estimates in Fernald (2012) are for the business sector rather than the nonfarm business sector.

However, the timing of these periods of capital deepening does not align straightforwardly with the timing of movements in labor's share. Although the 1990s witnessed significant growth in K/L , the period since 2000, during which labor's share has fallen substantially, has instead been marked by a considerable *slowdown* in growth in capital intensity.

For a given rate of labor-augmenting technical progress, then, the paths of real wages, productivity and the capital-labor ratio since 2000 diverge from that implied by an account of the decline in labor's share based on capital deepening. A potential reconciliation, of course, is that the rate of labor-augmenting technical progress fell since 2000. Two considerations add a note of caution to that interpretation. First, any such declines in the growth of A_L would have to be so precipitous that they more than offset the acceleration of growth in the effective capital-labor ratio k at the heart of the capital deepening story. Second, measures of investment-specific technological change have exhibited a similar slowdown since 2000. Figure 8 reveals that, although the initial acceleration in investment-specific technological change coincided with the onset of the decline in the labor share, the acceleration of the decline in the labor share over the last decade was not accompanied by an acceleration of investment-specific technological change. On the contrary, investment-specific technological change slowed down during that period.

The evidence suggests that the origins of the falling labor share in the United States do not dovetail easily with the neoclassical theory of distribution based on a simple aggregate production function. This has two possible implications: Either neoclassical theory is unable to provide a coherent account of the decline in labor's share, or a simple aggregate production function is too crude to capture the relevant economic forces. In the remainder of this section, we explore the latter by investigating the potential role of richer production structures and factor heterogeneity.

III.B. Skill-biased technological change and capital-skill complementarity

The simple aggregative benchmark analyzed above makes strong assumptions on the nature of the production process. A prominent example is its abstraction from heterogeneity in capital and labor, for example by vintage and by skill. There are empirical reasons to suspect that such heterogeneities have played an important role in recent decades in the U.S. economy. In particular, the substantial growth in wage inequality in the U.S. since the early 1980s has motivated two prominent and related hypotheses on the structure of technology.

The first is skill-biased technical change, the notion that technical progress particularly augments the productivity of high-skilled workers relative to the low skilled, yielding rising wage inequality (Bound and Johnson 1992; Katz and Murphy 1992). The second, capital-skill complementarity, explores the possibility that the elasticity of substitution between capital and skilled labor is less than that between capital and unskilled labor (Griliches 1969; Berman, Bound and Griliches 1994). In this case, declines in the relative price of new investment goods since the 1980s can stimulate capital deepening and thereby an increased skill premium due to capital-skill complementarity (Krusell, Ohanian, Rios Rull and Violante 2000).

In isolation, neither explanation provides an obvious account for a decline in the labor share. Skill-biased technical change can arise even in the presence of constant factor shares (for example, a Cobb-Douglas technology between capital and a constant-elasticity of substitution aggregate of labor by skill). Likewise, by operating through movements in the capital-labor ratio, the capital-skill complementarity story will affect the labor share for the same reasons noted in the subsection above—that is, if the elasticity of substitution between capital and labor (regardless of skill) differs from unity. Indeed, in their quantitative analysis of capital-skill complementarity, Krusell, Ohanian, Rios-Rull and Violante (2000) find very small effects on the labor share.

To see why, consider a constant-returns-to-scale aggregate production function with three factors, $y = F(k, s, u)$, where capital k , skilled labor s , and unskilled labor $u = 1 - s$ all are measured in units of effective labor. Mirroring equation (9), it is possible to show that labor's share will move according to the following relation:

$$d \ln \lambda = -(1 - \lambda) \left\{ [\omega_s \rho_{ks} + (1 - \omega_s) \rho_{ku}] d \ln k + \left[\frac{s}{1 - s} (1 - \omega_s) \rho_{ku} - \omega_s \rho_{ks} \right] d \ln s \right\}, \quad (10)$$

where $\rho_{ks} \equiv (\sigma_{ks} - 1)/\sigma_{ks}$ and $\rho_{ku} \equiv (\sigma_{ku} - 1)/\sigma_{ku}$ measure the degree of substitutability²⁷ between capital and labor by skill, and ω_s is the skill share of labor income. Equation (10) reiterates the point that the effect of capital deepening on labor's share in this more complicated environment mirrors those noted above in equation (9), where the overall degree of substitutability between capital and labor takes the form of a weighted average of ρ_{ks} and ρ_{ku} .

²⁷ The σ s here correspond to the so-called Hicks direct partial elasticities of substitution for this case of three inputs. While the concept of elasticity of substitution becomes more complex with more than two inputs (indeed, a variety of measures has been proposed), practical implementations of capital-skill complementarity often have used production functions for which these distinctions do not matter (see, for example, the two-level CES specification used by Krusell, Ohanian, Rios Rull and Violante 2000).

What is new is that changes in the skill mix of the labor force s can cause the labor share to move in the presence of capital-skill complementarity ($\rho_{ks} < \rho_{ku}$). Possible reasons for such a change include growth in the supply of highly educated workers, and skill-biased technical change. Moreover, the magnitudes of such effects of “skill deepening” could be substantial. For example, Acemoglu and Autor (2011) present estimates of relative weekly wages and supply for college versus high school educated workers using March Current Population Survey data for the years 1963 to 2008. These suggest that the share of college-educated workers (in efficiency units) s rose by around 75 log points over this period, and average values for the skill share of labor income and relative supply of around $\omega_s \approx 0.5$ and $s/(1-s) \approx 1$. Krusell, Ohanian, Rios Rull and Violante (2000) provide estimates of capital-skill complementarity parameters of $\rho_{ks} \approx -0.5$ and $\rho_{ku} \approx 0.4$. Combining these with a reasonable value for labor’s share, $\lambda = 0.6$, suggests a contribution of skill deepening to a decline in λ on the order of $-0.4 \times (1 \times 0.5 \times 0.4 + 0.5 \times 0.5) \times 75 = -13.5$ log points, strikingly similar to the order of magnitude seen in the data.

Despite this, a difficulty with this line of reasoning relates to the timing of movements in the labor share. Figure 9 illustrates the point. Rises in the skill share of labor s have occurred continually since the 1960s in the United States. Thus, the implied labor share, which is computed using a rolling year-by-year implementation of equation (10), is also predicted to fall *continually* throughout the sample period. As we have emphasized before, the reductions in the empirical labor share in Figure 3 are concentrated after the 1980s, preceded by relative trend stability, in contrast to the predictions of Figure 9. Put another way, the United States has not experienced a sharp acceleration in the skill share of labor in recent decades in a manner consistent with the timing of movements in the labor share.

III.C. Taking stock

Our analysis of a range of neoclassical explanations suggests that they provide an incomplete account of the movements in the labor share observed in recent U.S. history. Simple aggregative stories based on a non-unit elasticity of substitution between capital and labor face difficulties in explaining the joint movements in real wages, productivity and capital-labor ratios over the last few decades. Increases in the skill composition of the labor force combined with capital-skill complementarity have difficulty in replicating the peculiar timing of declines in the labor share.

Of course, our analysis does not definitively rule out the possibility that some other neoclassical explanation is able to fit the facts. What it does suggest, however, is that the canonical production structures that have been highlighted by recent and past literature fall short when confronted with the empirical movements in the U.S. labor share, especially those that have emerged during the last decade.

With an eye to the results of the analyses that follow, we close this section with a discussion of a different channel, namely the potential role of international trade in shaping movements in the labor share. We continue the theme adopted throughout this section of viewing the economy through the lens of an aggregate production function. Here, however, we abstract from technical change and focus instead on the role of intermediate inputs imported from abroad. Thus, production is given by $y = F(k, m)$, where y , k and m respectively denote output, capital and imported inputs per unit of labor. Note that we are interested in labor's share of *domestic* income—that is, output net of payments to imported factors. Under constant returns, the effect of increased use of imported intermediates is given by

$$\frac{\partial \ln \lambda}{\partial \ln m} = -(1 - \lambda)\mu(\rho_{lm} - \rho_{km}), \quad (11)$$

where ρ_{lm} and ρ_{km} index the substitutability²⁸ of labor and capital with respect to imported inputs, and μ is the income share of m in gross output y . Taken from an aggregate perspective, then, increased import exposure will reduce labor's share if imported intermediates are more substitutable with labor than they are with capital, $\rho_{lm} > \rho_{km}$. Intuitively, imported intermediates will be expected to induce greater substitution away from labor than from capital in this case. Our examination of disaggregated data in the ensuing sections shows that there is something to that expectation. For now, though, we note that the timing of the recent expansion of trade with developing economies accords somewhat better with the timing of the trend decline in the labor share since the late 1980s relative to accounts based on capital deepening.²⁹ In what follows, we show that there is a parallel link between these two trends in more disaggregated data.

²⁸ As before, here $\rho_{ij} \equiv (\sigma_{ij} - 1)/\sigma_{ij}$, and the σ s are the Hicks direct partial elasticities of substitution.

²⁹ A recent aberration to the general trend of increased trade has been the Great Recession, which was accompanied by a significant decline in trade, the so-called “trade collapse.” Nevertheless, this collapse has been almost entirely reversed since the recession; import penetration in the U.S. economy remains significantly above its pre-1990s level.

IV. Exploiting cross-industry variation

Since the most common aggregate production function representations fail to match the recent path of the labor share, in this section we consider several explanations that exploit cross-industry variation in payroll share declines documented in section II. We first revisit the investment-specific technological change hypothesis from an industry perspective. Available data for the period 1987 to 2011 suggest that it is not the case that the industries that saw the smallest increases in equipment prices are also those with the highest payroll share declines. We then consider two alternative explanations for the declines in labor shares across industries: (i) declines in unionization, and (ii) increased import competition. Our data yield one robust correlation: that declines in payroll shares are more severe in industries that face larger increases in competitive pressures from imports.

Investment-specific technical change and variation in equipment price inflation. Our analysis in Section 3 suggests a weak link between the decline in the labor share and investment-specific technical change. However, that analysis is based on a simple aggregate production function framework and could potentially be too crude to capture some of the changes in the production structure.

In this section, we instead exploit cross-industry variation to explore whether those industries experiencing the biggest declines in their payroll shares are also those that tend to invest in types of equipment whose price increases more slowly. That is, if the decline in the payroll share is due to firms replacing workers with machines, then sectors that invest in types of equipment with slower price increases should see larger declines in their payroll shares.³⁰ This would result in a positive relationship between the rate of equipment price increases and the change in the payroll share across industries.

Figure 10 plots this relationship using our data. On the horizontal axis is the average annual percent change in the price of equipment and software an industry invests in for 1987 through 2011. These data are taken from the Fixed Asset Tables made available by the Bureau of Economic Analysis. On the vertical axis is the percentage point change in an industry's payroll share between 1987 and 2011. We use the same sample of 60 NAICS industries that underlies the post-1987

³⁰ This is most often stated in the relative price of equipment (relative to output). Here we focus on the equipment price itself. Unreported results for the relative equipment price are qualitatively similar to the ones presented here.

results in Section II. We exclude oil and gas extraction from the sample, since movements in its payroll share are driven almost completely by changes in the price of oil. Industries in the scatterplot are split up into manufacturing, the squares, and nonmanufacturing, the circles. The size of each of the markers is proportional to the average value-added share of the corresponding industry over the period 1987 to 2011. The dashed line is the weighted-least squares regression line, where the industries are weighted by their value added shares.

A first interesting observation from Figure 10 is that declines in both payroll shares and equipment prices were commonplace from 1987 to 2011. Around three quarters of industries experienced a decline in their payroll share; around half of them experienced a decline in the equipment investment price. However, substantial declines in equipment price have not always been associated with payroll share declines. For example, the three industries with the biggest investment price declines saw either a small decline, or even an increase, in their payroll share.

Consequently, as Figure 10 shows, there is in fact a weak *negative* relationship between the change in equipment prices and payroll shares across industries. This is the opposite of what one would expect if capital deepening due to the decline in price of equipment were the driving force of the decline in the payroll share. This evidence strengthens the findings of our aggregate analysis and suggests that a shift to more capital-intensive production methods to exploit declines in equipment prices has not been a major factor in the evolution of the payroll share over the past 25 years.

Deunionization. An important change to the institutional structure of the U.S. labor market has been the decline in unionization. The fraction of workers covered by a union or employee association contract in the private nonagricultural sector has almost halved from 14.6 percent in 1987 to 7.7 percent in 2011.³¹ The bargaining power of unions tends to increase workers' share of the surplus generated in the production process. For example, Hirsch (2012) estimates that union wage premiums in the private sector were around 25 percent in 1984.³² Given the existence of a substantial union wage premium, a large decline in union membership might be expected to result in a decline in the aggregate labor share. Such a decline would in turn be expected to be concentrated in the industries with the largest declines in union coverage.

³¹ The unionization data we use are described in Hirsch and Macpherson (2003).

³² Hirsch (2012) finds that while there has been a modest downward trend in union wage premium in the private sector, it still remains at around 20 percent in 2010.

To explore this possibility, we make use of data on the change union coverage rates by industry between 1987 and 2011. Figure 11 plots these data against the change in payroll shares. From the figure it can be seen that, although there is a positive correlation between the change in unionization and the change in payroll shares across industries, the relationship is weak. The weighted least squares regression (Column II in Table 4) indicates that cross-industry variation in changes in unionization rates explains less than 5 percent of the variation in changes in payroll shares across industries. Though statistically insignificant, at 0.22, the point estimate of the effect of deunionization on the payroll share implies that in this simple regression deunionization can explain 1.5 percentage points of the drop in the aggregate payroll share.

Globalization: the role of rising imports. There has been a substantial increase in imports in the United States in the last few decades.³³ Part of this rise is a consequence of *vertical specialization*, which occurs when countries specialize in particular stages of a good's production sequence, rather than in the production of the entire good, as discussed by Hummels, Ishii and Yi (2001) and Yi (2003). In this subsection, we investigate the potential effects of such *offshoring* on recent trends in the U.S. labor share.

At the close of section III we noted that increased import penetration would be expected to depress labor's share of domestic income if imported intermediates are more substitutable with labor than with capital from the perspective of an aggregate production technology. From a microeconomic viewpoint, there are at least two channels by which such an effect could emerge. The first is reminiscent of the capital-labor substitutability hypothesis outlined in section III, except that movements in capital intensity are envisioned as the outcome of international trade.³⁴ Standard comparative advantage arguments imply that, when firms segment the production process, they will shift constituent parts to foreign countries with comparative advantage in the production of that specified part. In many industrialized countries, including the United States, this shift has taken the form of moving the production of labor-intensive intermediate and final goods to labor-abundant countries with lower labor costs (see, for example, Feenstra and Hanson 1996a, 1999). By offshoring the more labor-intensive part of U.S. production, the remaining production in the U.S.

³³ The ratio of imports to GDP went up from 10 percent in 1993 to around 16 percent in 2010.

³⁴ Of course, it is also possible that the processes of technical change emphasized in section III and increased penetration of imports of intermediate goods emphasized here are linked.

economy would be expected to become more capital intensive. If, in addition, capital is more than unit-elastic with respect to labor, Hicks' (1932) result will imply that the U.S. labor share will fall.

It is worth noting that this account is plausibly consistent with declines in labor shares in both source countries, such as the U.S., but also in the *destination* countries to which production is offshored, such as China. This is important if, as Karabarbounis and Neiman (2013) suggest, declines in labor's share have occurred globally. In particular, it is possible that offshored production processes that are labor intensive by U.S. standards also are capital intensive relative to existing production in China. And, indeed, there are precedents for this type of result in prior literature. Feenstra and Hanson (1996b), for example, note the importance of increased capital flows from the United States to Mexico in reconciling wage movements associated with offshoring between these two economies.

A difficulty with this first trade-based explanation, however, is that it faces the same challenges associated with other accounts based on capital deepening, namely that the timing of such deepening in the United States does not accord well with the timing of the decline in labor's share (see Figure 8). The second possible trade-based channel that we highlight does not rely on shifts in capital intensity, but instead takes as its point of departure the abundant evidence for inter-industry wage differentials documented by Dickens and Katz (1987) and Krueger and Summers (1988), among others. Viewed from this perspective, industry differences in labor shares have a more prosaic interpretation: Industries with high shares are those for which labor is expensive relative to its average productivity. It is natural to posit that these elements of the production structure would be particularly sensitive to import exposure, leading to a lowering of labor's share in such sectors.³⁵

In what follows, we assess the role of increased competition that U.S. workers are facing due to rising imports, leaving aside the important question of identifying the underlying economic channel for future work. To do so, we compute what we call *import exposure* for each of the industries that we consider. To do so, we pose the following question: If the United States were to produce domestically all the goods that it imports,³⁶ how much additional value added would each industry

³⁵ Again, note that this explanation is amenable to the possibility that labor's share will fall also in offshoring destination economies. If, for example, offshored wages are determined in part by low market wages in the destination country, and average labor productivity is determined in part by superior technologies from developed source economies, it is natural to imagine that offshored production units will have labor shares that are low relative to elsewhere in the destination economy.

³⁶ At current factor prices and productivity levels.

have to produce? For example, if all U.S. imports of clothes were produced domestically, how much would value added increase in sectors like retail, textile manufacturing, and so on.

To be able to calculate this measure of import exposure we use the annual input-output matrices that are available for the years 1993 to 2010 from the BLS. Import exposure is expressed as the percentage increase in value added needed to satisfy U.S. final demand if the U.S. would produce all its imports domestically.³⁷

Figure 12 plots the relationship between changes in import exposure and changes in the payroll shares for our sample of industries. The figure indicates that import exposure increased for almost all the industries; the aggregate pattern is not driven by a few sectors. It also reveals that those industries that faced the biggest rises in their import exposure also tended to experience the largest declines in their payroll share. Column III in Table 4 shows that the weighted least squares regression behind the plotted regression line is significant at the 0.1 percent level. It implies that the variation in the change in import exposure explains 22 percent of the cross-industry variation in payroll share changes.

The final column of Table 4 contains the results of a multivariate regression that jointly evaluates the quantitative importance of the three potential explanations considered in this section—investment-specific technical change, deunionization, and globalization. This shows that the effect of increased import competition on payroll shares across industries is robust to the inclusion of the other two explanatory variables. Moreover, once the change in import exposure is included the coefficient on the change in union coverage rates becomes effectively zero. This suggests that those sectors where deunionization was most prevalent are also sectors that saw the biggest increase in import exposure. The residual variation in the change in unionization rates does not explain much of the decline in the labor share.

Of course, this analysis is based on a set of simple cross-industry regressions and is thus subject to many caveats. Still, our results indicate a cross-industry link between the increases in import exposure and the decline at in the labor share. While this result cannot be interpreted as causal, it is worth noting that the statistical relationship between import exposure and payroll shares across

³⁷ In terms of input-output terminology, our measure of import exposure is the percentage difference between total domestic requirements and total requirements for an industry. One important aspect this measure does not capture is international roundaboutness of trade, where U.S. exports are used as intermediate abroad for the production of goods that then end up being imported into the U.S..

industries is of a magnitude large enough to account for a substantial fraction of the aggregate trend decline in the labor share. In particular, aggregating the results of the weighted least squares regression reported in column III of Table 4 across industries suggests that increases in import exposure of U.S. businesses can account for 3.3 percentage points of the 3.9 percentage point decline in the U.S. payroll share over the past quarter century.³⁸ This suggests that a particularly fruitful avenue for future research will be to delve further into the causal channels that underlie this statistical relationship, in particular the possibility that the decline in the U.S. labor share was driven by the offshoring of the labor-intensive component of the U.S. supply chain.

V. Practical implications and future developments

In the previous four sections we uncovered the main facts behind and explanations for the decline in the U.S. labor share. We conclude by considering whether these facts and explanations warrant a major rethink of how the labor share is used in empirical macroeconomics. In addition, with these factors and explanations in mind, we discuss what might drive movements in the labor share in the foreseeable future.

V.A. Reconsidering the use and interpretation of the labor share

There are three important uses of the labor share in applied macroeconomics. We discuss each of these separately in this subsection.

Factor income share The most straightforward interpretation of the labor share is as a measure of the distribution of income between the productive factors of capital and labor. It is in this context that Kaldor (1957, 1961) suggests that constancy of the share is a reasonable approximation to the data that should be replicated by models of economic growth. In practice, the relative constancy of the labor share has often been interpreted as an endorsement of the use of a Cobb-Douglas representation of aggregate technology in growth models. Of course, the use of a Cobb-Douglas production function does not explain the relative constancy of factor shares, it just postulates it.

The problem with gauging whether movements in the labor share dispel the stylized fact of constancy is that, as Solow (1958, p. 618) put it, “The literature does not abound in precise

³⁸ This 3.3 percentage point estimate is computed from the point estimate 0.87 times the value-added weighted average increase in import exposure across industries.

definitions, but obviously literal constancy is not in question.” So, are the recent movements in the U.S. labor share of an order of magnitude that justifies a reconsideration of the assumption of approximate constancy of factor shares? The answer to this question depends on the context to which the labor share as a measure of factor shares is applied.

The most common application of factor income shares is in growth and development accounting, in particular in the measurement of total factor productivity (TFP). In its simplest form, TFP growth is the difference between output growth, $d \ln Y$, and a factor-income-share-weighted average of the growth rates of labor, $d \ln L$, and capital, $d \ln K$.

Suppose that the true labor share is λ but that the measured share used for the growth-accounting exercise is $\lambda^* = \lambda + \varepsilon$. In that case, measured TFP growth, $d \ln TFP^*$, equals

$$d \ln TFP^* = d \ln Y - \lambda^* d \ln L - (1 - \lambda^*) d \ln K = d \ln TFP + \varepsilon(d \ln K - d \ln L). \quad (12)$$

This differs from actual TFP growth, $d \ln TFP$, by an amount equal to the measurement error, ε , times the growth rate of the capital-labor ratio. The latter is depicted in Figure 8 and averages 2.27 percent over the postwar period. Thus, if one were to use an estimate of the labor share consistently 5 percentage points higher than the actual share—around the order of magnitude of the recent decline—then, on average, one would obtain an estimate of annual TFP growth just 0.11 percentage points higher than its true counterpart.

To put the assumption of constancy of the labor share in this context into perspective, we recalculate Fernald’s (2012) measure of TFP growth with the labor share held constant at its postwar average. Figure 13 plots both the time-varying and constant labor shares used in these calculations, as well as the two implied measures of TFP growth constructed using these two labor share measures. As can be seen from the figure, to all intents and purposes both measures of TFP growth are the same.³⁹

Therefore, in spite of the recent decline in the labor share, it is important to bear in mind that there remain prominent applications, such as growth and development accounting, for which these movements do not make an appreciable difference.

The business cycle, markups, and resource slack There are many theoretical and empirical applications that emphasize and exploit the cyclical fluctuations in the labor share. As King and

³⁹ In fact, Basu, Fernald and Kimball (2006) use a constant labor share for their analysis.

Watson (2012) point out in the context of the New-Keynesian Phillips curve, such theories and applications hinge on the assumption that the labor share is mean stationary. In fact, mean stationarity is another way of interpreting “approximate constancy” of the labor share. King and Watson go on to show that the recent decline in the labor share has rendered it non-stationary. As we have shown, this is true not only for the headline measure but also for its two alternatives as well as for the payroll share. This means that uses of the labor share as a cyclical measure will have to address which part of its movements are cyclical and which parts reflect long-run trend movements.

The most common of such uses is that of the labor share as a proxy for wage markups and resource slack. Since the work of Galí and Gertler (1999) and Sbordone (2002), the headline labor share measure has often been used as a proxy for resource slack and marginal cost fluctuations for the estimation of the New-Keynesian Phillips curve.⁴⁰ Besides the difficulties with the recent time-series properties of the labor share in this context, a more fundamental conceptual issue is at stake. Many parts of compensation captured in the numerator of the labor share arguably are not part of firms’ marginal costs, in the sense that they do not represent costs that vary with labor input. For example, it could be argued that imputed hourly compensation of the self-employed is not part of firms’ direct payments to labor and thus not of its marginal costs. This suggests that the payroll share may be a better proxy for wage markups than the labor share. However, even some components of payroll compensation, such as stock options, are not linked directly to the quantity of labor used by the firm. Hence, the labor share is not a satisfactory measure of the wedge between the marginal cost and marginal product of labor. For these reasons, it is important to look beyond the labor share for measures of the wage markup. One potential way forward might be to focus instead directly on wages, rather than compensation, and productivity.

Proxy for inequality A second interpretation of the labor share is not as a measure of the division of income between the productive factors of labor and capital, but instead between workers and capitalists. In line with this, International Labor Organization (2013) interprets the recent decline in the labor share as an indication of inequitable growth.

As has been widely documented, and reiterated in Figure 4, the period during which the U.S. labor share declined was accompanied by a substantial increase in income inequality. It is natural,

⁴⁰ For example, the Fed’s FRB/US model uses such a specification for its inflation forecast (Brayton and Tinsley, 1996, Table 6, page 22).

then, to conjecture that the two processes are linked. Two aspects of Figure 4 strike a note of caution, however. First, the increase in inequality within labor income dwarfs the movements in the labor share. A second, related, feature of Figure 4 is that the labor share has in fact been *propped up* by the labor income of very highly paid individuals in recent decades. That is, to some extent, labor's share has declined *despite* the increased earnings of the super-rich.

To give an example, executive pay is part of labor compensation and is thus included in the labor share. Thus, the \$13 million in compensation, \$7.2 million severance payment, shares worth \$3.56 million, and performance bonus of \$2.4 million, that one CEO got during his 10-month stint at a major tech company in 2010 and 2011 all contributed positively to the labor share during those two years. In this sense, then, the decline in the labor share is an inadequate measure of the increase in income inequality.

By the same token, a further implication of Figure 4 is that the measured decline in the labor share would be much larger if not for the gains of the top 1 percent of payroll and self-employment income. By 2010, the labor share of the bottom 99 percent of taxpayers had fallen to approximately 50 percent from just above 60 percent prior to the 1980s. Aggregate measures of the labor share miss this richer detail in the income structure, so that the decline in the aggregate share understates significantly the increase in income inequality in the U.S. over the last 25 years.

V.B. Drivers of future movements in the labor share

Our analysis of a range of factors behind and explanations for the recent decline in the labor share highlights that the decline of the labor share over the last 25 years is largely driven by U.S. producers facing increased import competition. Thus, if globalization continues during the next decades, then the labor share will most likely continue to decline, especially in sectors that face the largest increases in foreign competition.

However, it is important to realize that sector-specific developments, such as the tech bubble at the end of the 1990s, can have a surprisingly big effect on movements in the labor share. This suggests that it is worthwhile also to consider more sector-specific developments going forward.

In addition to developments in certain industries affecting the aggregate labor share, particular types of legislation are also likely to drive movements in the near future. The Affordable Care Act (ACA) is one example. Since health insurance benefits are part of payroll compensation, a change

in the prevalence of these benefits that is not offset by changes in wages and salaries will affect the payroll and labor shares. Moreover, since the legislation makes health insurance more accessible for the self-employed, it will likely result in an increase in the share of the labor force that is self-employed, thus changing the wedge between the labor share and the payroll share. This will affect the measured labor share. Exactly how will depend on the way the BLS decides to treat self-employment income in its measurement of the labor share going forward.

References

- Acemoglu, Daron, and David Autor (2011), “Skills, Tasks and Technologies: Implications for Employment and Earnings,” *Handbook of Labor Economics Volume 4*, Orley Ashenfelter and David E. Card (eds.), Amsterdam: Elsevier.
- Alvaredo, Facundo, Anthony B. Atkinson, Thomas Piketty and Emmanuel Saez (2013), *The World Top Incomes Database* [<http://topincomes.g-mond.parisschoolofeconomics.eu>], last accessed 07/29/2013]
- Antras, Pol (2004), “Is the U.S. Aggregate Production Function Cobb-Douglas? New Estimates of the Elasticity of Substitution,” *B.E. Journal of Macroeconomics* 4(1). Autor, David H., Lawrence F. Katz, and Melissa S. Kearney (2008), “Trends in U.S. Wage Inequality: Revising the Revisionists,” *Review of Economics and Statistics*, 90(2), 300–323.
- Basu, Susanto, John G. Fernald, and Miles S. Kimball (2006), “Are Technology Improvements Contractionary?” *American Economic Review*, 96, 1418-1448.
- Bentolila, Samuel and Gilles Saint-Paul (2003), “Explaining movements in the labor share,” *Contributions in Macroeconomics*. Volume 3, Issue 1, Pages –, ISSN (Online) 1534-6005.
- Blanchard, Olivier (1997), “The Medium Run,” *Brookings Papers on Economic Activity*, 89-158.
- Berman, Eli, John Bound, and Zvi Griliches (1994), “Changes in the Demand for Skilled Labor within U.S. Manufacturing: Evidence from the Annual Survey of Manufacturers,” *Quarterly Journal of Economics* 109(2): 367-397.
- Bosworth, Barry, and George L. Perry (1994), “Productivity and Real Wages: Is There a Puzzle?” *Brookings Papers on Economic Activity*, 1994-1, 317-344.
- Bound, John, and George Johnson (1992), “Changes in the Structure of Wages in the 1980s: An Evaluation of Alternative Explanations,” *American Economic Review*, 82(3): 371-392.
- Brayton, Flint, and Peter Tinsley (1996), “A Guide to FRB/US,” *Finance and Economics Discussion Series 1996-42*, Federal Reserve Board of Governors.
- Chirinko, Robert S. (2008), “ σ : The Long and Short of It,” *Journal of Macroeconomics* 30(2): 671–86.

- Dickens, William T. and Lawrence F. Katz (1987), "Interindustry Wage Differences and Industry Characteristics," in K. Lang and J. Leonard, eds. *Unemployment and the Structure of Labor Markets*. New York: Basil Blackwell.
- Estrada, Ángel, and Eva Valdeolivas (2012), "The Fall of the Labour Income Share in Advanced Economies," *Occasional Paper Series #1209*, Bank of Spain.
- Feenstra, Robert C. and Gordon H. Hanson (1996a), "Globalization, Outsourcing, and Wage Inequality," *American Economic Review*, 86 (2), 240-245.
- Feenstra, Robert C. and Gordon H. Hanson (1996b), "Foreign Investment, Outsourcing and Relative Wages," in Robert C. Feenstra, Gene M. Grossman and Douglas A. Irwin, eds., *The Political Economy of Trade Policy: Papers in Honor of Jagdish Bhagwati*. Cambridge: MIT Press.
- Feenstra, Robert C. and Gordon H. Hanson (1999), "The Impact of Outsourcing and High Technology Capital on Wages: Estimates for the United States, 1979-1990," *Quarterly Journal of Economics*, 114 (3), 907-40.
- Fernald, John G. (2012), "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity," *FRBSF Working Paper 2012-19*.
- Galí, Jordi, and Mark Gertler (1999), "Inflation Dynamics: A Structural Econometric Analysis," *Journal of Monetary Economics*, 44, 195-222.
- Gallaway, Lowell E. (1964), "The Theory of Relative Shares," *Quarterly Journal of Economics*, 78(4), 574-591.
- Gollin, Douglas (2002), "Getting Income Shares Right," *Journal of Political Economy*, 110, 458-474.
- Gomme, Paul, and Peter Rupert (2004), "Measuring Labor's Share of Income," *Policy Discussion Paper - No 7*, Federal Reserve Bank of Cleveland.
- Grant, Arthur (1963), "Issues in Distribution Theory: The Measurement of Labor's Relative Share, 1899-1929," *Review of Economics and Statistics*, 45(3), 273-279
- Greenwood, Jeremy, Zvi Hercowitz and Per Krusell (1997), "Long-Run Implications of Investment-Specific Technological Change," *American Economic Review* 87(3): 342-362.
- Griliches, Zvi (1969), "Capital-skill Complementarity," *Review of Economics and Statistics*, 51(4), 465-468.
- Guscina, Anastasia (2006), "Effects of Globalization on Labor's Share in National Income," IMF Working Paper, WP/06/294
- Harrison, Ann (2002), "Has Globalization Eroded Labor's Share? Some Cross-Country Evidence," mimeo, UC Berkeley.
- Hicks, John R. (1932), *The Theory of Wages*. London: Macmillan.
- Himmelberg, Charles P., James M. Mahoney, April Bang, and Brian Chernoff (2004), "Recent Revisions to Corporate Profits: What We Know and When We Knew It," *FRBNY Current Issues in Economics and Finance*, 10(3), March 2004.

- Hirsch, Barry T., and David A. Macpherson (2003), "Union Membership and Coverage Database from the Current Population Survey: Note," *Industrial and Labor Relations Review*, 56(2), pp. 349-54. [Data available at www.unionstats.com, last accessed 09/06/2013]
- Hirsch Barry T. (2012), "Unions, Dynamism, and Economic Performance," in *Research Handbook on the Economics of Labor and Employment Law*, Michael Wachter and Cynthia Estlund, eds., Edward Elgar Series of Research Handbooks in Law and Economics.
- Hummels, David, Jun Ishii and Kei-Mu Yi (2001), "The Nature and Growth of Vertical Specialization in World Trade," *Journal of International Economics*, 54(1), 75-96.
- Hurst, Erik, Geng Li and Benjamin Pugsley (2010), "Are Household Surveys Like Tax Forms: Evidence from Income Underreporting of the Self Employed," NBER Working Paper #16527.
- Hurst, Erik, Geng Li and Benjamin Pugsley (2012), "Are Household Surveys Like Tax Forms? Evidence from Income Underreporting of the Self-Employed," *Review of Economics and Statistics*, forthcoming.
- International Labor Organization (2013), *Global Wage Report 2012/13: Wages and Equitable Growth*, Geneva: International Labour Office.
- Johns, Andrew, and Joel Slemrod (2010), "The Distribution of Income Tax Noncompliance," *National Tax Journal*, 63, 397-418.
- Jorgenson, Dale W. (1963), "Capital Theory and Investment Behavior," *American Economic Review (Papers and Proceedings)*, 53(2), 247-259.
- Kaldor, Nicholas (1957), "A Model of Economic Growth," *Economic Journal*, 67, 591-624.
- Kaldor, Nicholas (1961), "Capital Accumulation and Economic Growth," in *Theory of Capital*, F.A. Lutz and D.C. Hague (eds), New York: St. Martin's Press.
- Karabarbounis, Loukas, and Brent Neiman (2013), "The Global Decline of the Labor Share," *NBER Working Paper No. 19136*.
- Katz, Lawrence F. and Kevin M. Murphy (1992), "Changes in Relative Wages, 1963-1987: Supply and Demand Factors," *Quarterly Journal of Economics* 107(1) 35-78.
- Keynes, John M. (1939), "Relative Movements of Real Wages and Output," *Economic Journal*, 49, 34-51.
- King, Robert G., and Mark W. Watson (2012), "Inflation and Unit Labor Cost," *Journal of Money Credit and Banking*, 44s, 111-149.
- Kravis, Irving B. (1959), "Relative Income Shares in Fact and Theory," *American Economic Review*, 49(5), pp. 917-949.
- Krueger, Alan B. (1999), "Measuring Labor's Share," *American Economic Review, Papers and Proceedings*, 89, 45-51.
- Krueger, Alan B. and Lawrence H. Summers (1988), "Efficiency Wages and the Inter-Industry Wage Structure," *Econometrica*, 56, 259-293.

- Krusell, Per, Lee E. Ohanian, José-Víctor Ríos Rull, and Giovanni L. Violante (2000), “Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis,” *Econometrica* 68(5), 1029-1053.
- Mehran, Hamid, and Joseph Tracy (2001), “The Effect of Employee Stock Options on the Evolution of Compensation in the 1990s,” *FRBNY Economic Policy Review*, Volume 7(3), 17-34.
- Moylan, Carol E. (2008), “Employee Stock Options and the National Economic Accounts,” *Survey of Current Business*, February 2008, 7-13.
- Piketty, Thomas, and Emmanuel Saez (2003), “Income Inequality in the United States: 1913-1998,” *Quarterly Journal of Economics*, 118(1), 1-39.
- Sbordone, Argia M. (2002), “Prices and Unit Labor Costs: A New Test of Price Stickiness,” *Journal of Monetary Economics*, 49, 265–292.
- Slemrod, Joel (2007), “Cheating Ourselves: The Economics of Tax Evasion,” *Journal of Economic Perspectives*, 21(1), Winter 2007, 25–48.
- Solow, Robert M. (1958), “A Skeptical Note on the Constancy of Relative Shares,” *American Economic Review*, 48(4), 618-631.
- Yi, Kei-Mu (2003), “Can Vertical Specialization Explain the Growth of World Trade?” *Journal of Political Economy*, 111 (1), 52-102.

Figure 1. Labor share, payroll share, and replicated labor share in U.S. nonfarm business sector.

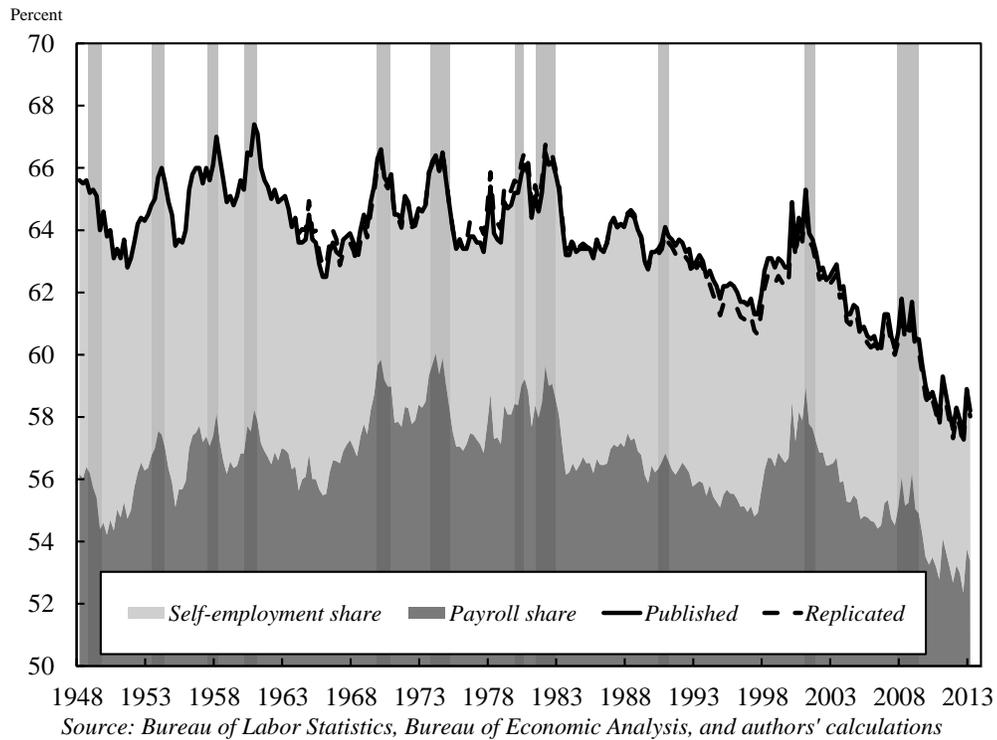


Figure 2. Composition of nonfarm business sector income.

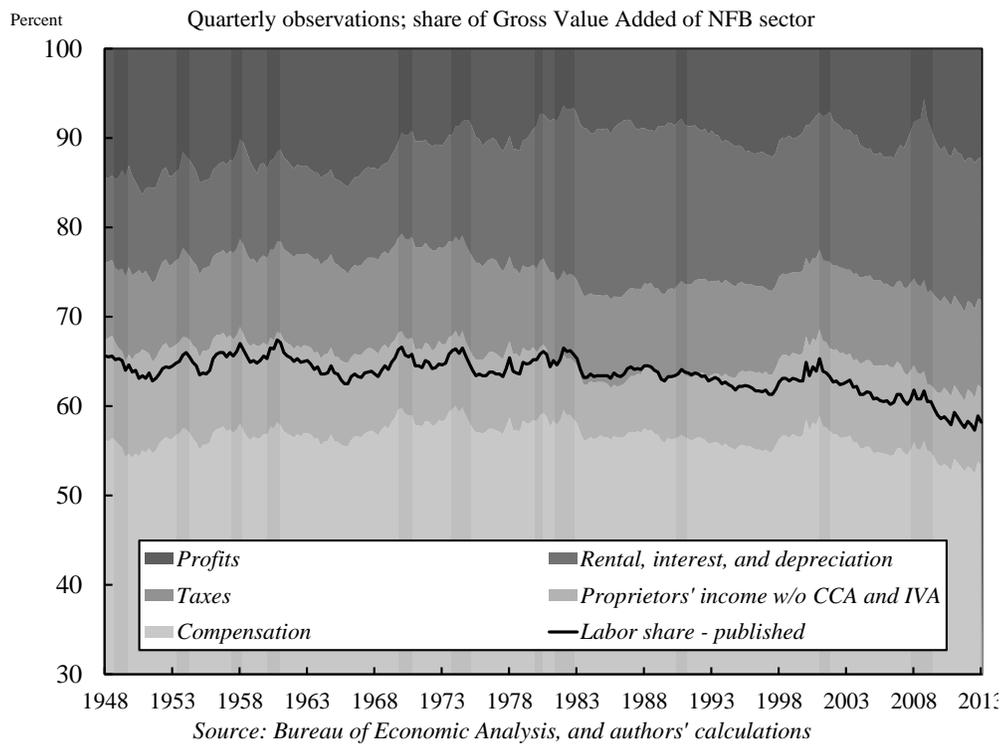


Figure 3. Alternative measures of labor share based on four estimates of self-employment labor income.

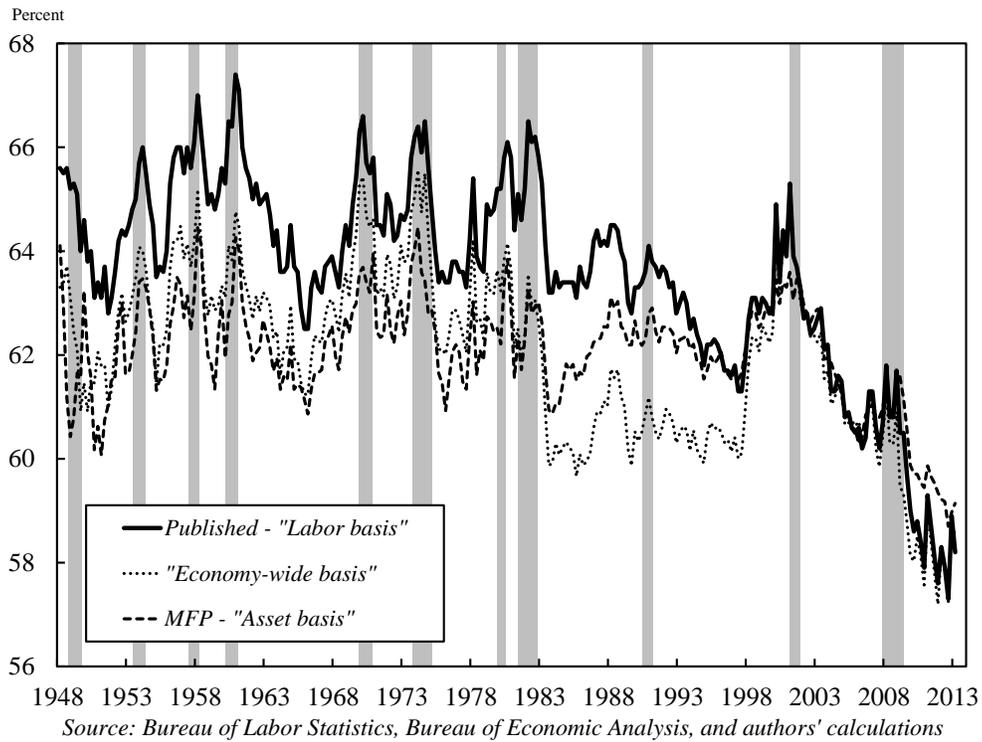


Figure 4. Nonfarm business labor share split up by income fractile.

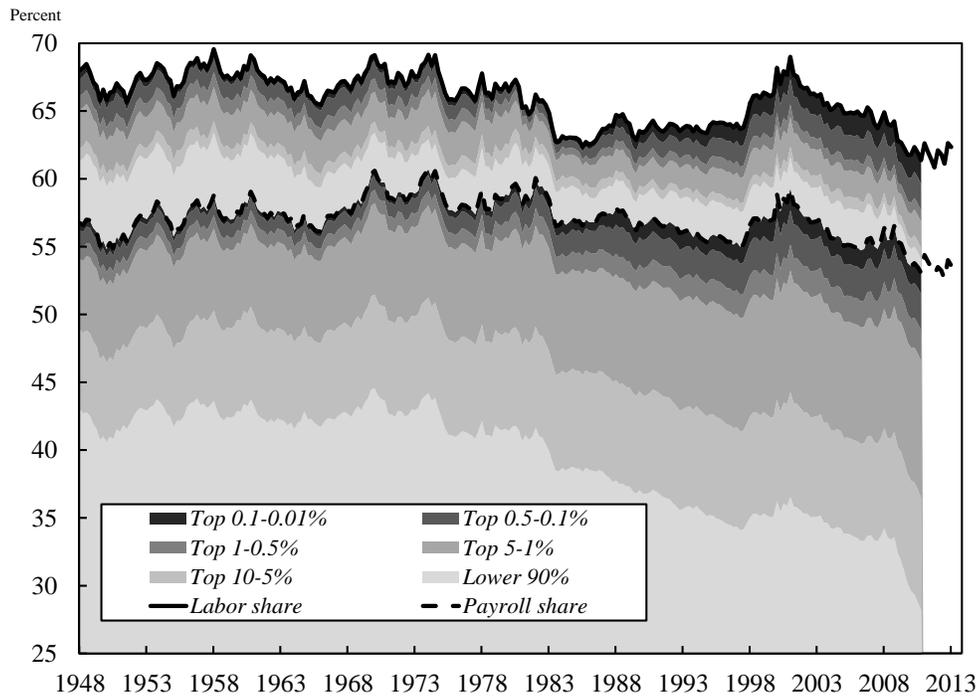


Figure 5. Cross-industry distribution of payroll share.

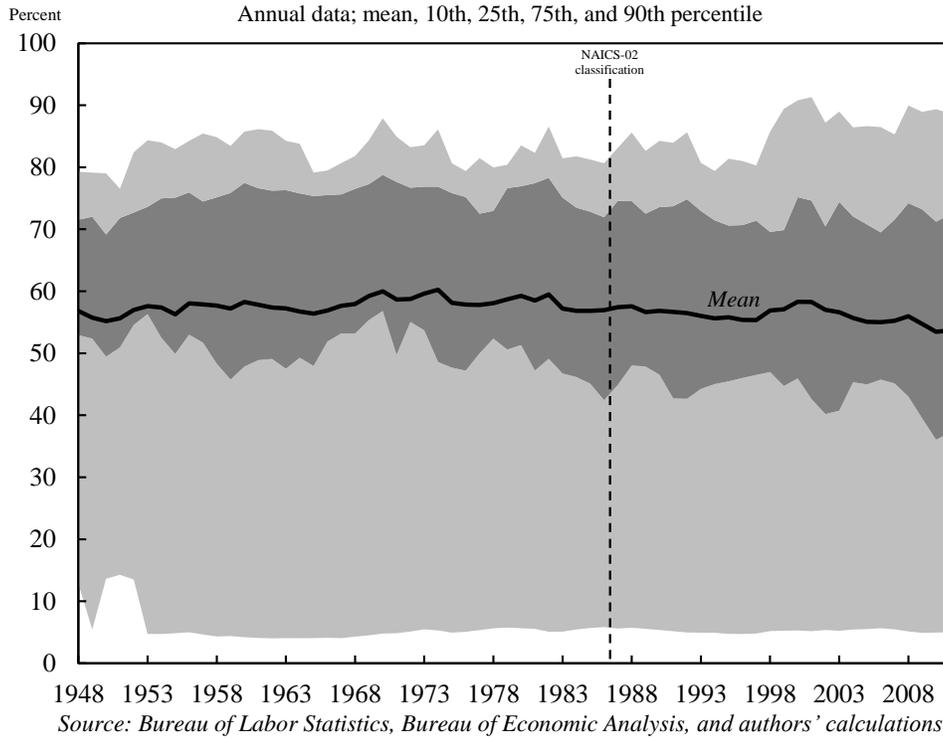


Figure 6. Contributions of changes in payroll shares in tech and investment banking to aggregate.

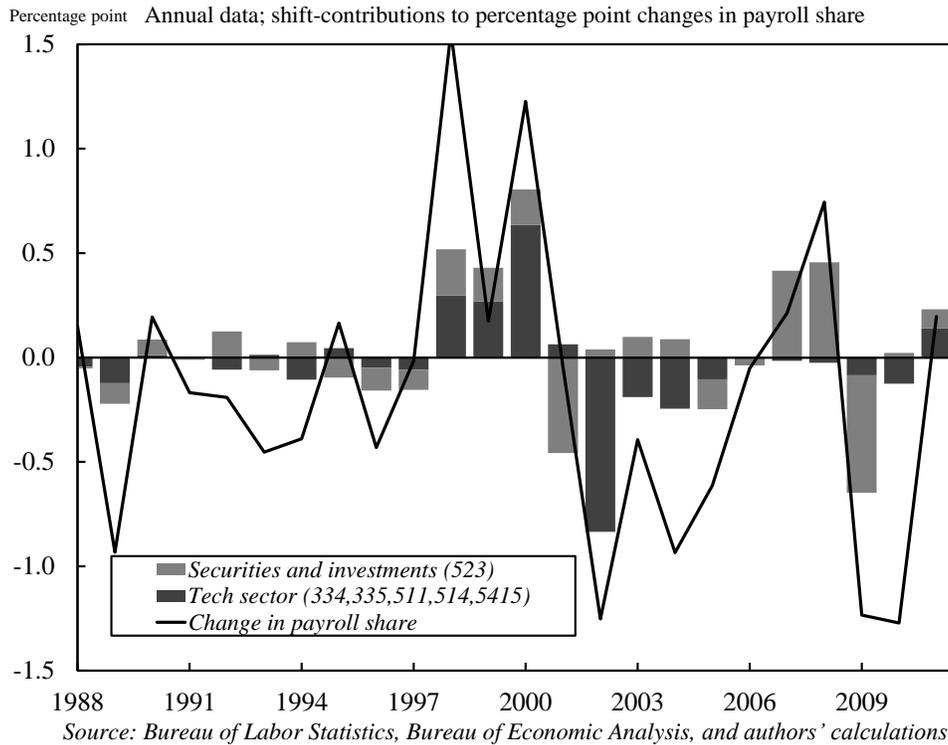


Figure 7. Growth of real hourly compensation and output per hour in the nonfarm business sector.

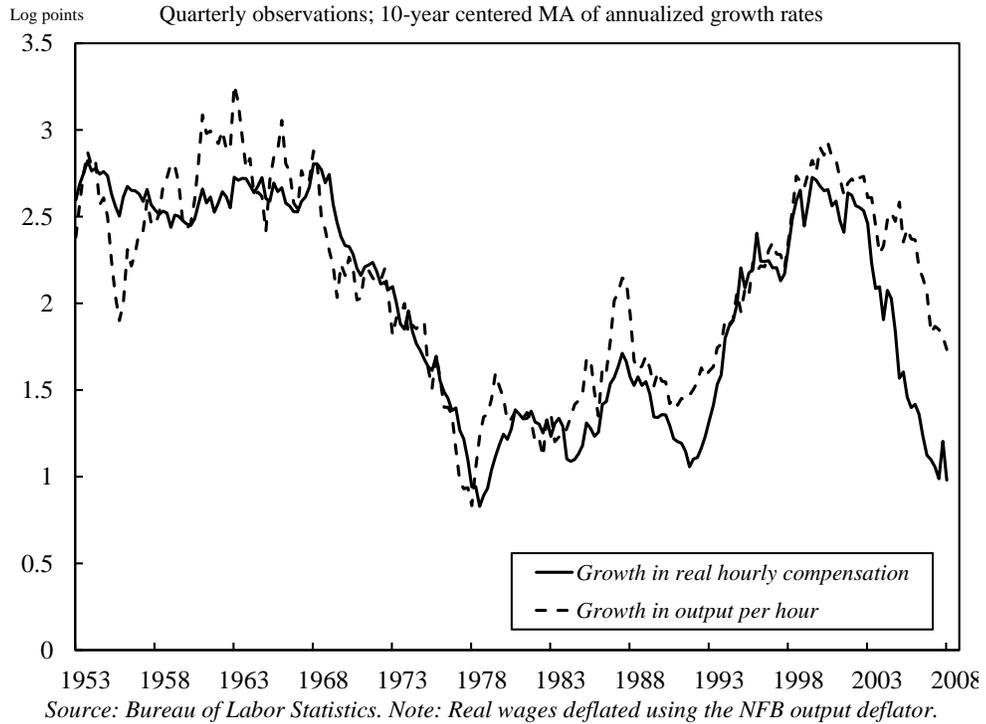
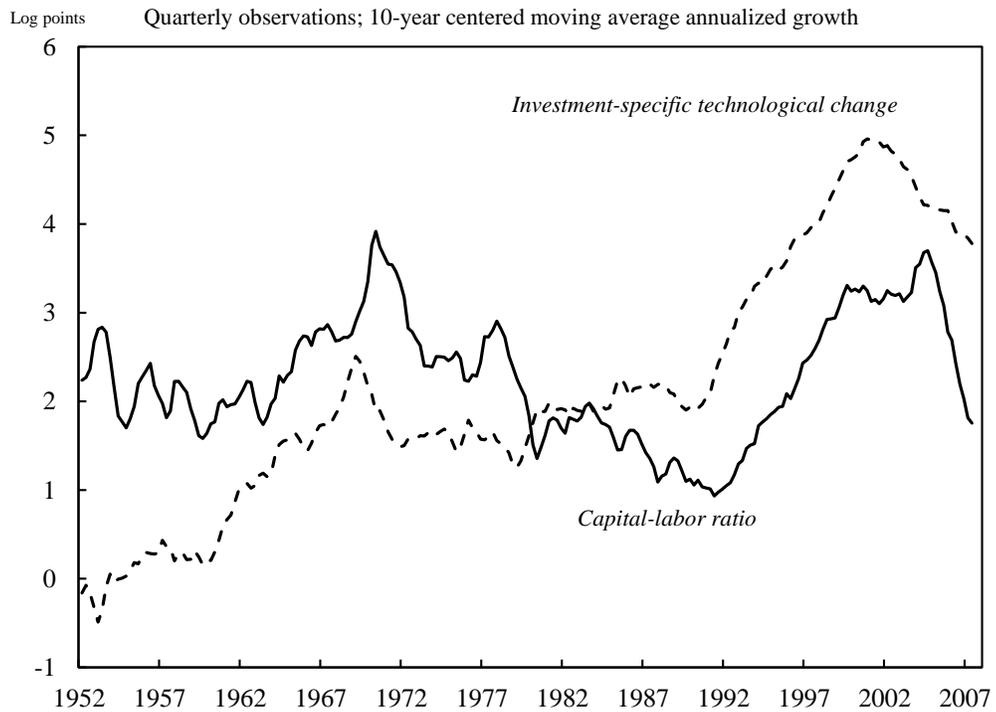
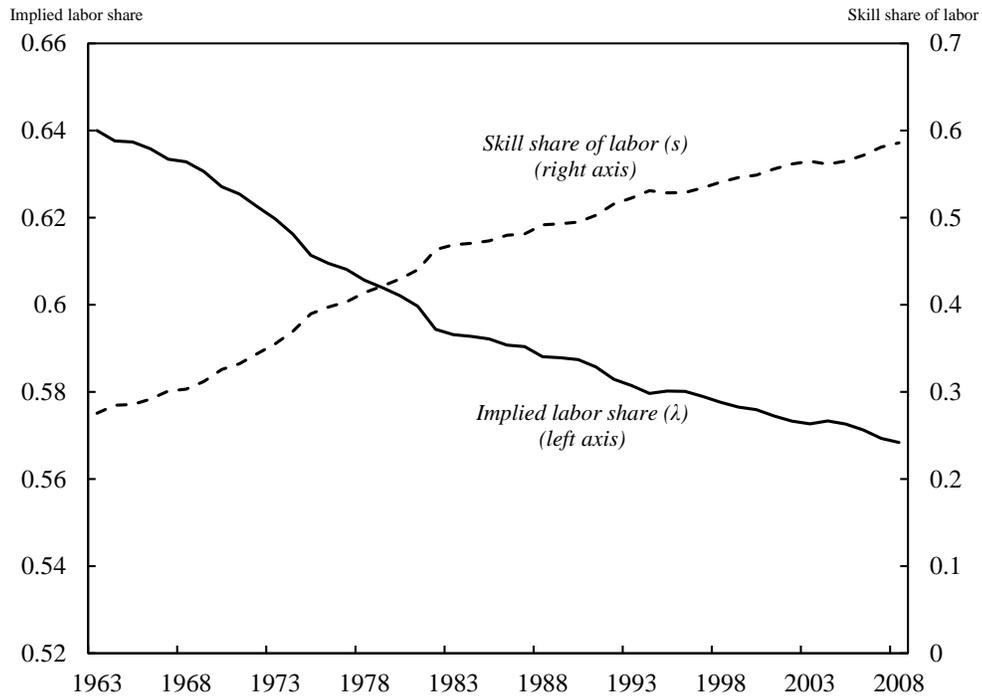


Figure 8. Growth of capital-labor ratio and investment-specific technological change in the business sector.



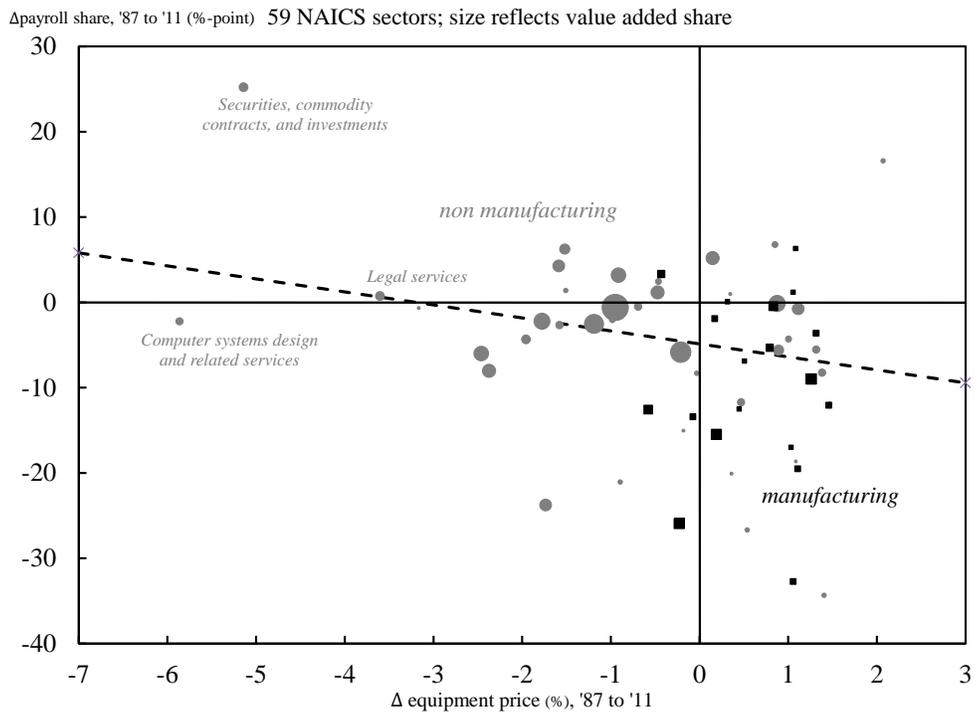
Source: Fernald (2012)

Figure 9. Theoretical path of labor share implied by skill deepening and capital-skill complementarity.



Source: Skill share of labor and implied path of labor share based on estimates of Acemoglu and Autor (2011) and Krusell et al. (2000).

Figure 10. Changes in payroll shares and equipment price inflation by industry.



Source: Bureau of Economic Analysis, Bureau of Labor Statistics, and authors' calculations

Figure 11. Deunionization and changes in payroll shares by industry.

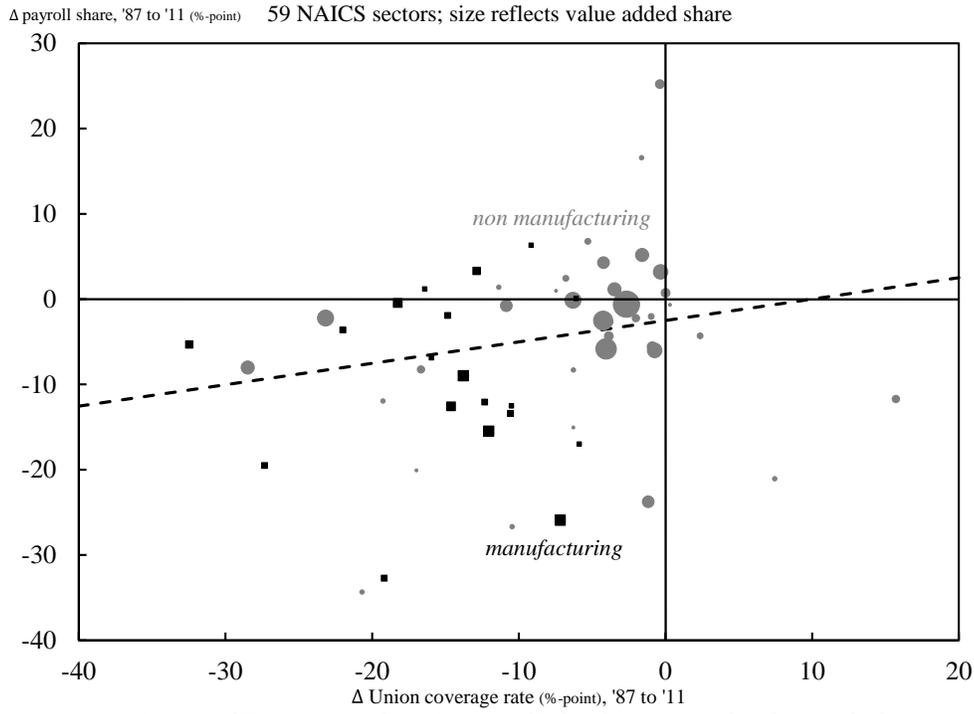


Figure 12. Import competition and changes in payroll shares by industry.

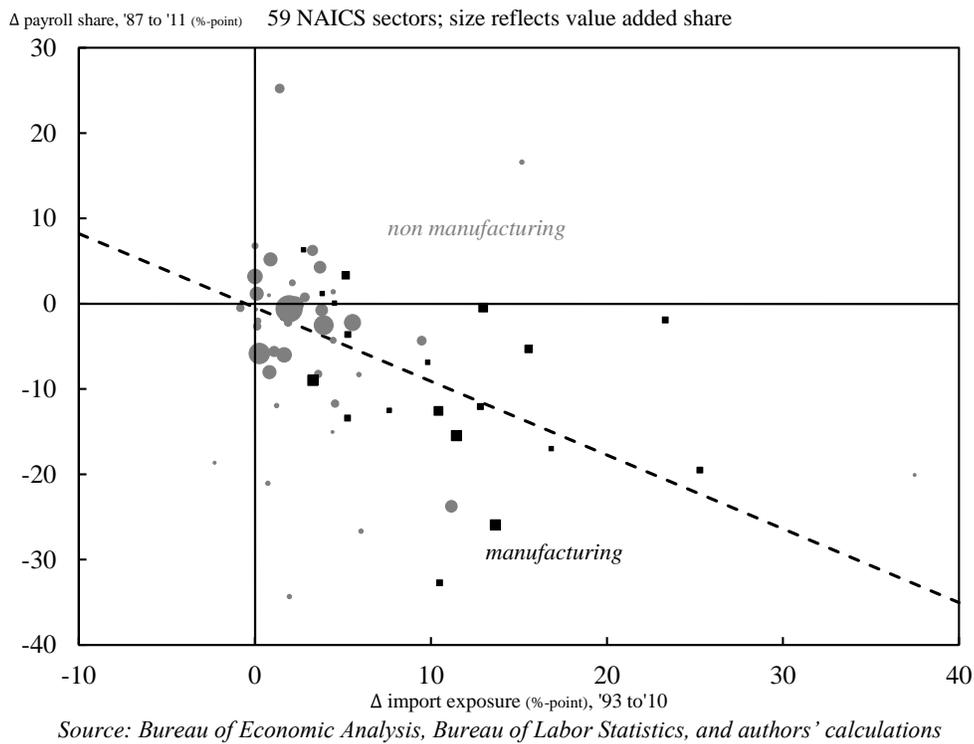
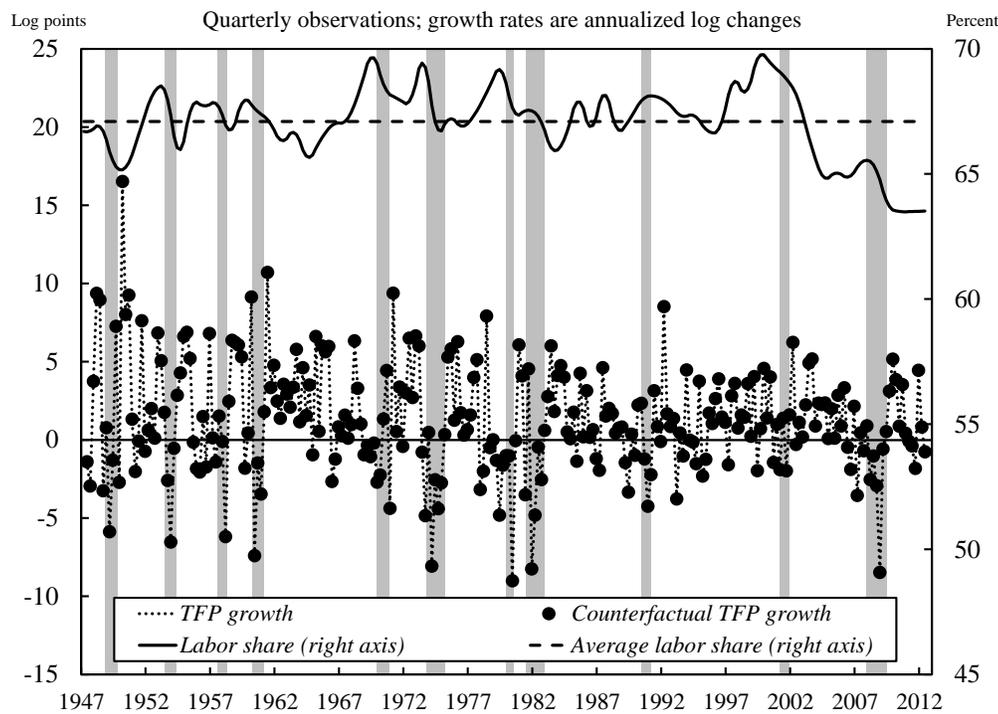


Figure 13. Measured TFP growth and counterfactual TFP growth under constant labor share.



Source: Fernald (2012) and authors' calculations. Results are for the business sector.

Table 1. Sources and magnitude of the decline in the labor share.

Measure	Mean	Mean '48 to '87	Mean '10 to '12	Change
<i>Composition of nonfarm business sector income</i>				
1. Compensation (payroll share)	56.5	57.1	53.3	-3.9
2. Proprietors' income w/o CCA and IVA	8.8	9.0	8.4	-0.6
3. Taxes	9.8	9.9	9.9	0.0
4. Rental, interest, and depreciation	13.9	12.4	16.4	3.9
5. Profits	11.0	11.6	12.1	0.5
<i>Alternative measures of the labor share</i>				
6. "Labor basis" – published	63.6	64.6	58.3	-6.4
7. "Asset basis" – MFP	62.6	62.9	59.3	-3.6
8. "Economy-wide basis"	62.0	62.7	58.1	-4.7

Source: Bureau of Economic Analysis, Bureau of Labor Statistics, and authors' calculations.

Note: Sample covers 1948Q1 to 2013Q1.

Table 2. Shift-share analysis of changes in payroll share by major sector, 1948 to 2011.**Panel A. SIC, 1948 to 1987**

Major Sector	Share of value added			Payroll share			Shift-share analysis	
	1948	1987	Change	1948	1987	Change	Shift	Share
Nonfarm business sector	100	100	0	56.8	57.4	0.6	3.5	-2.7
Natural resources and mining	4.7	3.0	-1.7	40.6	36.8	-3.8	-0.1	-0.7
Construction	5.3	5.5	0.2	70.4	69.9	-0.5	0.0	0.1
Durable goods manufacturing	17.6	12.9	-4.7	76.6	78.5	1.9	0.3	-3.6
Non-durable goods manufacturing	15.3	7.6	-7.7	61.3	63.0	1.7	0.2	-4.7
Trade/Transportation and utilities	31.7	26.0	-5.7	56.9	59.2	2.3	0.8	-3.3
Information	4.1	5.0	0.9	66.1	51.9	-14.2	-0.6	0.5
Financial activities	12.6	20.5	7.9	20.1	26.3	6.3	1.0	1.8
Professional and business services	1.3	6.2	5.0	65.0	70.7	5.6	0.2	3.3
Education and health services	2.4	7.1	4.7	50.0	83.7	33.6	1.6	3.1
Leisure and hospitality	1.4	1.6	0.2	55.7	63.0	7.3	0.1	0.1
Other services	3.5	4.4	1.0	61.6	61.5	-0.1	0.0	0.6

Panel B. NAICS, 1987 to 2011

Major Sector	Share of value added			Payroll share			Shift-share analysis	
	1987	2011	Change	1987	2011	Change	Shift	Share
Nonfarm business sector	100	100	0	57.4	53.7	-3.8	-3.6	-0.1
Natural resources and mining	2.2	2.5	0.3	42.4	32.2	-10.1	-0.2	0.1
Construction	5.2	4.1	-1.1	72.6	72.4	-0.2	0.0	-0.8
Durable goods manufacturing	12.3	7.0	-5.2	78.7	68.4	-10.3	-1.0	-3.8
Non-durable goods manufacturing	8.2	6.3	-1.9	62.3	40.4	-21.8	-1.6	-1.0
Trade/Transportation and utilities	22.6	19.3	-3.4	58.2	53.8	-4.4	-0.9	-1.9
Information	4.9	5.0	0.1	47.8	42.5	-5.3	-0.3	0.0
Financial activities	21.2	23.6	2.5	24.8	24.9	0.1	0.0	0.6
Professional and business services	9.6	14.6	5.0	70.9	72.5	1.6	0.2	3.6
Education and health services	6.9	10.1	3.2	85.0	86.3	1.2	0.1	2.7
Leisure and hospitality	3.8	4.6	0.8	65.7	62.1	-3.6	-0.1	0.5
Other services	3.0	2.9	-0.2	64.9	70.0	5.2	0.2	-0.1

Source: Bureau of Economic Analysis, Bureau of Labor Statistics, and authors' calculations.

Table 3. Variance decomposition of annual changes in the aggregate payroll share, 1948 to 2011.

Major Sector	SIC '49 to '87			NAICS '87 to '11		
	Total	Shift	Share	Total	Shift	Share
Nonfarm business sector	100	102	-2	100	109	-9
Natural resources and mining	4	-1	5	1	-1	2
Construction	10	5	5	13	7	6
Durable goods manufacturing	33	36	-3	16	19	-3
Non-durable goods manufacturing	9	11	-2	3	8	-5
Trade/Transportation and utilities	24	28	-5	15	20	-5
Information	5	5	0	10	17	-7
Financial activities	3	10	-7	11	15	-4
Professional and business services	5	3	2	29	13	17
Education and health services	3	2	2	-5	3	-8
Leisure and hospitality	1	1	0	4	5	-1
Other services	2	2	1	2	3	0

Source: Bureau of Economic Analysis, Bureau of Labor Statistics, and authors' calculations.
Note: Percentage point contributions to the variance of changes in the aggregate payroll share.

Table 4. Regressions to explain cross-industry variation in changes in payroll shares, 1987 to 2011.

Dependent variable: Δ payroll share (%-point), '87 to '11	I	II	III	IV	V
Change in equipment price (%), '87 to '11	-1.52 (0.78)			-1.03 (0.72)	-0.90 (0.74)
Change in unionization coverage rate (%-point), '87 to '11		0.22 (0.14)			0.01 (0.13)
Change in import exposure (%-point), '93 to '10			-0.87*** (0.21)	-0.81*** (0.22)	-0.82** (0.24)
R^2	0.06	0.04	0.22	0.25	0.25
No. of observations	59	56	59	59	56

*Note: Significance levels * $p < 5\%$, ** $p < 1\%$, *** $p < 0.1\%$. All results are obtained using weighted least squares with an industry's average value-added share over 1987 to 2011 as the weight.*