# Growing Productivity without Growing Wages: The Micro-Level Anatomy of the Aggregate Labor Share Decline*

Matthias Kehrig† and Nicolas Vincent ‡

First version: November 2015, this version: March, 2018
ERID Working Paper Number 244

**Abstract**

The aggregate labor share in U.S. manufacturing declined dramatically over the last three decades: Since the mid-1980's, the compensation for labor declined from 67% to 47% of value added which is unseen in any other sector of the U.S. economy. The labor share of the typical U.S. manufacturing plants, in contrast, rose by over 5 percentage points. We reconcile these two facts by documenting (1) an important reallocation of production towards "hyper-productive plants" and (2) a polarization of the labor share distribution over time. Yet, while we observe that low-labor-share plants have become relatively larger, we also find that the labor share of "hyper-productive plants" remains strikingly transient. This can be explained by a sharp increase over time between the size of an establishment and the likelihood that it features a low labor share in the future. These "tilted lottery" micro dynamics have important implications for our understanding of the drivers of the aggregate labor share decline.

Keywords: Labor Share, Productivity, Firm Size Distribution, Organization of Markets.
JEL classification: E2, L1, L2, L6, O4.

*We would like to thank Nick Bloom, John Cochrane, Allan Collard-Wexler, Steve Davis, Chad Jones, Pete Klenow, Luigi Pistaferri and our discussant Kirill Shakhnov as well as seminar participants at Stanford, Penn State, Mannheim, the 2015 meeting of German Economists Abroad, the Annual Meeting of the SED, the Barcelona Summer Forum, the UniCredit Alumni Workshop the Société Canadienne de Science Économique and the Cambridge-INET conference "Firms in Macroeconomics," for helpful comments. All errors are our own. Financial support from the Fritz Thyssen Stiftung and the Fondation HEC Montréal is gratefully acknowledged. Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed.
The most recent version can be downloaded here: https://ssrn.com/abstract=2943059.

†Duke University, Department of Economics, Box 90097, Durham, NC 27708, matthias.kehrig@gmail.com.

‡HEC Montréal, Department of Applied Economics, 3000 Chemin de la Côte-Sainte-Catherine, Montréal, Quebec H3T 2A7, Canada, nicolas.vincent@hec.ca.

# 1 Introduction

Several recent studies have documented a decline of the aggregate labor share, the portion of gross domestic product paid out in compensation for labor. This finding is very important for a number of reasons: it contradicts one of the stylized facts of Kaldor (1961) which have become foundational for theories of economic growth. It is further at odds with a key building block of standard macroeconomic models, the Cobb-Douglas production function. Lastly, it seems to suggest that an economy's value added gets distributed less to those who produce that value added and more to those that own the means of production.

A growing literature is studying potential causes for this aggregate labor share decline and comes to conflicting conclusions: Elsby et al. (2013) identify offshoring of labor-intensive activities as the main driver of the aggregate labor share decline. Karabarbounis and Neiman (2014a) attribute the labor share decline world-wide to increasingly higher efficiency of new capital. Eden and Gaggl (2018) find that the rise of information and communication capital has replaced routine labor and thus lowered that portion of the labor share. Koh et al. (2016) claim that a higher intangible capital intensity is responsible for the decline in the labor share. Corrado et al. (2005, 2009) made headways in measuring intangible capital and propose that its deepening explains most of labor productivity growth, a concept inversely related to the labor share. Rognlie (2015) identifies a low-frequency fall and rise of the capital share and shows that its recent increase is mostly due to the housing sector.

The common theme across these studies is that they are concerned with finding an aggregate explanation for a labor share pattern they document at the aggregate or sectoral level. What is unknown at this point are the micro-level dynamics that underpin the aggregate labor share decline. Understanding these dynamics will in turn help us better grasp the forces that underlie the fall in the labor share. The present paper fills this gap and uses confidential data from the U.S. Census of Manufactures to study the micro-level anatomy of the aggregate labor share decline. Our focus is on U.S. manufacturing, a sector for which detailed plant-level data is available and that sidesteps or minimizes some of the measurement challenges highlighted above, such as the roles of self-employed income, intangibles or housing. We confirm that the labor share in the manufacturing sector declines by almost 5 percentage points (ppt) per decade between 1967 and 2007. This, however, hides a striking fact: Alongside this aggregate decline, the median establishment actually saw an *increase* in its labor share, by about 1.4 ppt per decade. In fact, this upward trend is present for the vast majority of manufacturing establishments. These two facts are depicted in Figure 1.

To reconcile the opposing trends of the labor share at the aggregate and establishment levels, we document the role of two related forces using a decomposition exercise. First, we show that there has been a dramatic reallocation of production from high- and mid-labor-share establishments towards low-labor-share establishments over this period. We label these latter establishments "hyperproductive plants." Second, we observe a fattening of the tails of the distribution of establishment-level labor shares over time. Per se, this polarization of labor shares should not necessarily have an impact on the aggregate trend. However, because the hyperproductive plants capture a larger and

larger portion of aggregate manufacturing value added, the fact that they are able to lower their labor share over time implies that this widening of the distribution also pushes down the aggregate. We also find that most of the downward adjustment in the manufacturing labor share has been happening in the years following recessions. This finding is consistent with the evidence on employment from various authors who note that the disappearance of routine jobs is an important factor behind the recent jobless recoveries experienced in the U.S. and elsewhere (see Acemoglu and Autor (2011); Jaimovich and Siu (2015)).

Figure 1: The changing distributions of labor shares and value added

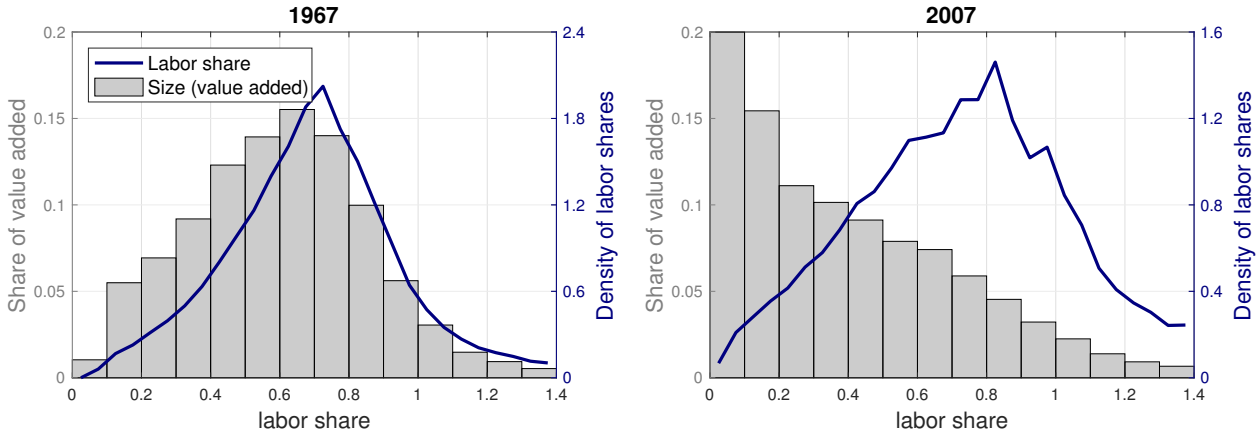


*Note:* The cross-establishment distribution of labor shares (solid blue lines) shows no significant locational shift of establishment-level labor shares from 1967 to 2007; the fattening of tails indicates a polarization of labor shares. The distribution of economic activity (value added shares in grey bars), in contrast, dramatically shifts towards low-labor-share establishments. This reallocation of value added is principally responsible for the aggregate labor share decline.

Taken together, these two interlinked forces account for almost all of the change in the evolution of the aggregate labor share since the early 1980s. We attribute a third of the downward pressure on the aggregate labor share to reallocation of economic activity from relatively high-labor-share to relatively low-labor-share establishments, while half of the downward pressure is driven by those establishments that simultaneously lower their labor share *and* grow disproportionately. Although this latter channel accounts for half of the observed aggregate labor share decline, it is accounted for by only the very small group of "hyperproductive plants." The remaining contributions, which are due to entry and exit as well as mergers and acquisitions, are very limited and add the remaining sixth of the forces pushing the aggregate labor share downward.

In the last part of the paper, we focus our attention on the "hyperproductive" plants that account for most of the decline in the labor share of the manufacturing sector. In order to achieve an average labor share that is more than 2.5 times smaller than that of their peers by 2007, we find that these hyperproductive plants relied almost solely on growing their value added while keeping their wage and employment levels in line with those of other establishments.

**Literature review** A burgeoning literature has documented and come up with different explanations for the labor share decline. One set of explanations involves technical change. Karabarbounis and Neiman (2014a) have put forward the notion that technical change embodied in new equipment capital has displaced labor and lowered the labor share. Eden and Gaggl (2018) refine this theory by focusing on information and communication technology capital. Koh et al. (2016) emphasize the rise of intangible capital such as intellectual property products, research and development and knowledge capital in the production function of developed economies. Alvarez-Cuadrado et al. (2015) show that industry-specificities in technological change and the elasticity of substitution between capital and labor matters for the dynamics of industry-level factor shares. A common ingredient in the argument of these papers is that the elasticity of substitution between equipment or intangible capital and (routine) labor has to be greater than unity. This has been criticized by Lawrence (2015) and Oberfield and Raval (2014) who present evidence that the elasticity of substitution is in fact smaller than unity, at various levels of aggregation.

Alternatively, Elsby et al. (2013) advocate the role of offshoring as the main driver of the labor share decline. In related work, Boehm et al. (2015) present establishment-level evidence that outsourcing did cut U.S. manufacturing employment while raising profits per worker of surviving production units. Glover and Short (2016) find the age composition of the work force has shifted towards workers that are less capable of extracting their marginal product of labor as a wage.

Hartman-Glaser et al. (2017) study Compustat data and find a similar dichotomy between the aggregate and average capital share that we find in labor share data. They explain the rise in the aggregate capital share through increasingly risky firm productivity. In their model, more volatile productivity implies that the firm owner can ask for a larger insurance premium, raising in turn the capital share. Even though Davis et al. (2006) have shown that establishment-level outcomes have become less volatile, Kehrig (2011) has shown that the productivity dispersion across establishments has increased significantly. This evidence can be reconciled with Davis et al. (2006) if spreads in the fixed productivity component of new establishments widen while the volatility of the variable components may well have declined. From the perspective of individual workers this widening would also pose an increased risk requiring more ex ante insurance.

Lastly, an emerging strand of the labor share literature emphasizes the role of rising concentration and markups: Autor et al. (2017a) present some industry-level evidence on firm concentration shares which is consistent with our finding that a small fraction of "hyperproductive plants" are mainly responsible for the aggregate labor share decline. Grullon et al. (2016) use firm-level data from Compustat to document that most U.S. industries became more concentrated over time, with the "winning firm" making large profits and realizing outstanding stock returns as well as more profitable mergers and acquisitions. Barkai (2017) shows that markups have grown over time, lowering both the labor and capital shares.

Issues related to the measurement of the labor share abound: Elsby et al. (2013) refine the imputation of the labor portion of noncorporate income which mitigates the labor share decline a bit. Bridgman (2014) claims that the rise of less durable capital such as computers and software

means that a larger share of value added is spent on replacing depreciated capital. Karabarbounis and Neiman (2014b) explore that issue using world-wide data and show that the potential of higher depreciation to explain the labor share decline is limited: broad trends in the *gross* and *net* labor shares are in fact quite similar.

# 2 The dynamics of the U.S. labor share

## 2.1 Data sources and measurement

We use both industry-level data from the Bureau of Labor Statistics (BLS) and confidential establishment-level data from the Census of Manufactures. The BLS data come from the annual "KLEMS Multifactor Productivity Tables by Industry" for both Manufacturing and Nonmanufacturing Industries and span the period from 1948-2014. We use the SIC based tables until 1987 and then switch to the NAICS based tables from 1987 onwards, adjusting the SIC-based time series so that the SIC and NAICS based times series coincide in 1987.

The U.S. Census Bureau collects data on all manufacturing establishments within the economic census, which is taken every five years from 1967 until 2007. We drop all observations which are administrative records or which are not part of the "tabbed sample" which make up the official tabulations published by Census.

In either dataset, the labor share $\lambda_t$ in a given industry and year $t$ is defined as

$$\lambda_t = \frac{w_t L_t}{Y_t} \tag{1}$$

where $w_t L_t$ denotes aggregate labor costs and $Y_t$ aggregate value added produced in the manufacturing sector at time $t$ gross of depreciation and taxes.

In the BLS data, labor costs comprise employee compensation (wages, salaries and supplements) as well as a portion of non-corporate income.[1] We compute value added as the value of production minus the costs for materials, energy inputs and purchased services.

In the Census data, we define the following items as labor costs: salaries and wages (item `SW`), involuntary labor costs (item `VLC`) such as unemployment insurance or social security contributions netted out from wages and voluntary labor costs (item `ILC`) such as health, retirement and other benefits paid to employees. Value added in the Census data is measured as sales less inventory investment for final and work-in-progress goods, resales[2], material inputs and energy expenditures.

[1] The "Technical Information About the BLS Multifactor Productivity Measures" (September 2007) states the assumptions involved in allocating non-corporate income to labor and capital costs in each year: "Initially self-employed persons and unpaid family workers are assumed to receive the same hourly compensation as employees and the rate of return to non-corporate capital is assumed to be the same as in the corporate sector. Based on these assumptions, the resultant income of proprietors is adjusted to match actual proprietors income reported in the GPO data by scaling proportionately the hourly compensation of the self-employed and the noncorporate rate of return. This treats any apparent excess or deficiency in noncorporate income neutrally with respect to labor and capital." (p. 9).

[2] This means we consider the value added by an establishment's production activities, not its trading activities.

Unlike in the BLS data, purchased services are not reported in the Census data. To account for that, we reduce establishment-level value added by the industry-year-specific ratio of purchased services to value added computed from the BLS data. In addition, we drop all observations outside the one percentiles to avoid outliers driving result. This means we also drop observations with a negative value added (and thus labor share). We do that in order to have one consistent Census sample in our analyses throughout the paper, some of which would not be meaningful because they involve observations with negative labor shares.

Both the BLS and Census measures lack any non-monetary compensation or ownership rights which have monetary value to an employee. Stock options, for example, are counted as labor income for tax purposes once a manager exercises the option but not at the point in time when the manager acquires the option. Ongoing research in finance is concerned with the rising share of deferred compensation in total labor compensation. This could potentially mitigate the aggregate labor share declines in both the BLS and our Census measure.

Compared to the aggregate BLS labor share, our aggregate labor share measure based on the micro-level Census data will be lower for three reasons. First, we do not include non-corporate (self-employed) compensation as part of labor compensation, so our numerator will be lower. Second, we do not consider establishments with negative value added, so our denominator will be greater. Third, our way to make up for the missing purchased services will likely leave value added higher as well, again making our denominator greater. These three factors imply that the aggregate labor share in manufacturing computed from the Census data is about eight percentage points smaller than in the BLS data.

## 2.2 The labor shares falls in most U.S. sectors

Before we delve into the plant-level manufacturing data, we first provide evidence on the evolution of the labor share across main U.S. sectors. For each industry, we compute the labor share as in equation (1), using sectoral data from the BLS as described in the previous section. Table 1 presents the share of value added of each sector in 2007 alongside its labor share decline over the 1987-2014 period.

The results indicate that while the drop in labor share is particularly pronounced for Manufacturing as well as other sectors such as Mining or Wholesale Trade, the overall trend is broad-based. Only a few sectors, such as Education or Arts & Recreation experienced a noticeable increase. This confirms us in our focus on the manufacturing sector: it is not an outlier, yet the fact that it experienced a clear and pronounced decline of its labor share makes us more confident that we can identify significant dynamics from the micro-level data. Even though it lost some share in the aggregate economy, it still is one of the main sectors and thus important enough to warrant attention. In addition, manufacturing is less prone to some of the issues highlighted in the literature, such as the role of self-employment income, intangibles or housing capital in driving the decline in the labor share.

Table 1: Labor share decline by sector

| Sector | Labor Share decline (in percentage points) | Share of value added 1987 | 2014 |
|---|---|---|---|
| Agriculture | +0.0 | 2.4 | 1.7 |
| Mining | −6.6 | 1.8 | 3.1 |
| Utilities | +0.7 | 4.1 | 2.7 |
| Construction | −1.2 | 5.4 | 4.9 |
| Manufacturing | −6.6 | 22.8 | 14.3 |
| Non Durables | −8.8 | 8.9 | 6.6 |
| Durables | −4.2 | 13.9 | 7.7 |
| Trade | −3.5 | 13.3 | 12.3 |
| Wholesale Trade | −4.1 | 5.8 | 6.3 |
| Retail Trade | −2.4 | 7.4 | 6.0 |
| Transport & Warehousing | −2.8 | 3.7 | 3.5 |
| Information | −0.5 | 4.7 | 5.7 |
| Finance, Insurance, Real Estate | −0.6 | 12.1 | 13.6 |
| Services | −0.3 | 19.9 | 26.7 |
| Health | −0.4 | 4.1 | 5.6 |
| Other | −1.0 | 5.8 | 6.1 |

*Note:* Authors calculations of sectoral labor share declines (in percentage points) per decade based on BLS data from the KLEMS Multifactor Productivity Tables by Industry, data range 1987-2014. Shares in nominal aggregate value added are in percent.
"Other" collects Education, Arts & Recreation, Hotels & Restaurants and Other Services.

## 2.3 The labor share in U.S. manufacturing

From now on we will focus on the labor share dynamics in the U.S. manufacturing sector. We choose to focus on that sector for two reasons. First, as shown above and also highlighted by Elsby et al. (2013), manufacturing is one of the sectors where the labor share decline is most pronounced, making it a good starting point to study the macro and micro dynamics of the labor share decline. Second, micro-level data from the U.S. Census Bureau are available down to the level of the individual establishment and for a long time period. This allows us to contrast the dynamics of the labor share both before and after the start of its secular decline, around the early 1980s. In our analysis, we will use the quinquennial Census of Manufactures which started in 1963 and were conducted every five years from 1967 onwards. This gives us a history of labor share dynamics that is longer than the ones available in other sectors, which the U.S. Census Bureau only started to sample in the 1980s.

**Manufacturing labor share declines strongly** As a first exercise, we confirm that the aggregate labor share in the U.S. manufacturing sector as measured in the Census establishment-level data is consistent with the labor share in the industry-level BLS data. To that end, we compute the aggregate labor share in the Census data by aggregating labor costs and value added across all plants in a given year to compute the numerator and denominator of equation (1). In Figure 2 we compare the aggregate manufacturing labor shares in both the BLS and the Census data from 1967 onwards.

Figure 2: The aggregate labor share in U.S. manufacturing

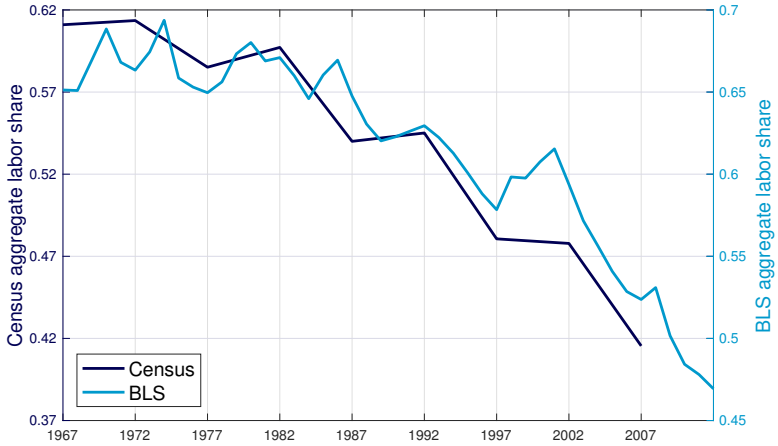


Figure 2 shows that the BLS labor share is about 8 ppt higher in general because both labor compensation and value added are measured differently in the BLS and Census data (see above). We also drop observations with negative value added in the Census data, thus lowering the labor share. Yet, the two labor series line up well in terms of level and, more importantly, trends. While the original work by Karabarbounis and Neiman (2014a) documents a 1.4 ppt decline per

decade in the global corporate sector, the labor share in manufacturing declines by a stunning 4.9 ppt per decade over our sample period. The vast majority of this decline occurred since the mid 1980s: Up to 1982, the manufacturing labor share fell by only a meager 0.9 ppt per decade while it dropped by 7.3 ppt per decade since the 1982 Census. Importantly, since we consider data from the producer side and focus on the manufacturing sector, our analysis is unlikely to be impacted by the measurement problems present in household-level data. For example, Elsby et al. (2013) argue that self-employment income was an important contributor to these trends. Conversely, Rognlie (2015) documents that income from housing alone was responsible for the labor share dynamics computed from household-side surveys, and Eden and Gaggl (2018) document a similar pattern for residential capital income in more aggregate income and product accounts. Our analysis, focusing on manufacturing plant level data, circumvents these issues.

**Labor share adjustment over the business cycle** Next, we show that there exist significant differences in the extent of the adjustment of the labor share over the various phases of the business cycle. Ríos-Rull and Santaeulàlia-Llopis (2010) noted that the labor share is coutercyclcial which may result from wages or labor that are relatively sluggish compared to output. This can be explained in a model of technology shocks with search frictions, labor adjustment costs and/or sticky wages due to Nash bargaining. We add to this cyclical pattern our new observation that the secular labor share decline occurs in expansions rather then recessions. Based on the annual BLS data for the manufacturing sector, we find that prior to the 1980s, the aggregate labor share generally increased by 2-3 ppt in a recession and declined by the same amount in the subsequent recovery. The medium slump in U.S. manufacturing in 1986 initially follows this pattern with a 2-percentage-point rise, but the subsequent recovery sees the size of the labor share drop heavily, by about 5 ppt. Every business cycle since has shown similar dynamics: a modest rise of the labor share in the recession, followed by a large decline thereafter. This pattern was particularly dramatic between the 2001 and 2009 recessions, with a decline of 9 ppt. This evidence adds to the job polarization literature. Jaimovich and Siu (2015), for example, who show the anatomy of jobless recoveries over the last 30 years: despite recoveries in aggregate output, employment picks up only slowly following the 1991, 2001 and 2009 recessions because the middle-occupation jobs eliminated in the previous recessions do not come back.

**Different declines in production and non-production compensation** The labor costs used in the numerator of the labor share contain various components. In the Census data, we can distinguish between production worker wages, salaries for non-production workers as well as ancillary labor costs. Production worker wages include the wage bill of all employees engaged in the core manufacturing activities such as fabricating, processing, assembling, inspecting, receiving, packing, warehousing, maintenance, repair, janitorial and guard services and record keeping. Salaries of non-production workers refer instead to the compensation of all employees above line-supervisor level. It comprises executive, purchasing, professional and technical sales, logistics, advertising, credit, clerical and routine office functions. Finally, the ancillary labor costs comprise legally required

labor costs (such as social security tax, unemployment tax, workmen's compensation insurance and state disability insurance pension plans) as well as voluntary labor costs (such as health benefits, life insurance premiums, supplemental unemployment compensation and deferred profit sharing plans).

We investigate whethere these three components declined symmetrically. This question is important as some theories of the labor share decline such as deunionization or the automation of routine jobs would be expected to have a disproportionately large impact on the wages of production workers, while affecting to a lesser degree the two other components. Other theories such as a change in the competitive landscape would likely have a more symmetric effect on all three labor share components.

We study this question by decomposing the aggregate labor share into these three components. The results are shown in Figure 3 and Table 2:

$$\lambda_t = \underbrace{\frac{w_t^{pw} L_t^{pw}}{Y_t}}_{\text{Wage bill}} + \underbrace{\frac{w_t^{npw} L_t^{npw}}{Y_t}}_{\text{Salaries}} + \underbrace{\frac{w_t^{ben} L_t^{ben}}{Y_t}}_{\text{Ancill. labor costs}} .$$

We find that the compensation of production workers declines secularly, by about 4.6 ppt per

Figure 3: Dynamics of labor share components

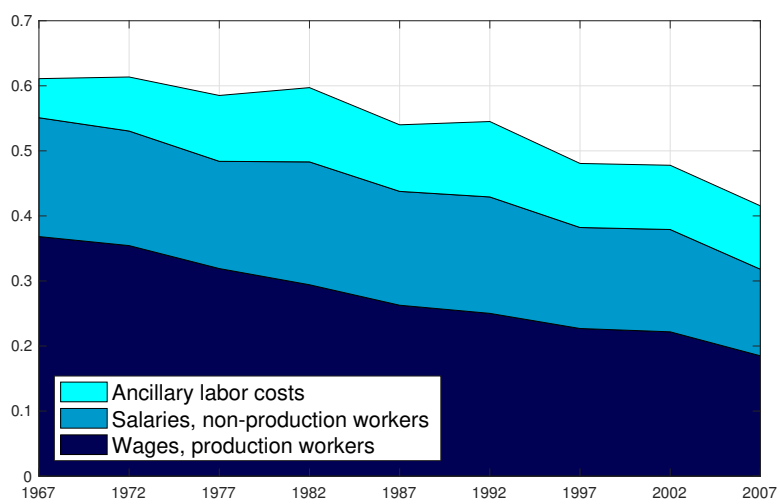


decade, mirroring the average rate of decline of the overall labor share. However, while the aggregate labor share stays roughly constant until the early 1980s, the compensation of production workers declines steadily since the beginning of our dataset in the late 1960s. In fact, once the downward trend in the overall labor share starts in the early 1980s, the compensation decline for production workers slows down slightly. All in all, had the production-worker labor share not declined at all, the aggregate labor share would have stayed more or less constant (-0.3 ppt per decade).

The compensation for non-production labor, in contrast, is steady at first and then starts to

Table 2: Dynamics of labor share components per decade (percentage point change)

| Component | 1967-2007 | 1967-1982 | 1982-2007 |
|---|---|---|---|
| | (percentage point changes) | | |
| Aggrgate labor share | −4.9 | −0.9 | −7.3 |
| | | | |
| Production worker wages | −4.6 | −4.9 | −4.4 |
| Non-production worker salaries | −1.2 | +0.4 | −2.2 |
| Ancillary labor costs | +0.9 | +3.6 | −0.7 |

decline after 1982, but not as strongly as that of production labor. If the compensation for non-production labor had stayed constant rather than declining at 1.2 ppt per decade, the aggregate labor share would have only declined by 3.7 ppt per decade instead of 4.9 ppt. Ancillary labor costs display the opposite pattern: they push the aggregate labor share up by almost one percentage point per decade. In the early decades of our data, the increase in the ancillary labor costs and salaries offset the decline in production worker wages thus leaving the aggregate labor share constant until 1982. After that, the ancillary labor costs decline only slightly. Had they not dampened the overall decline of labor compensation, the aggregate labor share decline would have been stronger at 5.8 ppt per decade instead of the observed 4.9 ppt decline.

## 2.4 The labor share decline across industries and regions

**Industry factors** The observed aggregate labor share decline could stem from specific industries which experience a decline while others experience only a small or no change in their labor share. Likewise, the aggregate labor share could fall because economic activity in terms of value added gets reallocated to relatively low labor-share industries.

To test for such compositional effects, we decompose the aggregate labor share decline into a component within and across industries using equation (2):

$$\Delta\lambda_t = \underbrace{\sum_j \Delta\lambda_{jt}\omega_{jt-1}}_{\text{Within adjustment}} \underbrace{+\sum_j \lambda_{jt-1}\Delta\omega_{jt}}_{\text{Between reallocation}} \underbrace{+\sum_j \Delta\lambda_{jt}\Delta\omega_{jt}}_{\text{Residual}}. \tag{2}$$

Table 3 displays the results from the within-between industry decomposition. It shows that most of the labor share decline from 1967-2007 stems from within-industry adjustment. Defining an industry at the 3-digit NAICS level, 3.3 ppts of the 4.9 ppt decline is due to within-industry adjustment, while between-industry reallocation only account for 0.7 ppts. The residual interaction term can be interpreted as either adjustment of relatively expanding industries or reallocation directed to industries that lower their labor share.

Table 3: Within- vs. between-industry factors in the labor share decline

| Portions of labor share change | 1967-2007 | 1967-1982 | 1982-2007 |
|---|---|---|---|
| | (percentage point changes) | | |
| Aggregate labor share change | −4.9 | −0.9 | −7.3 |
| | | | |
| *NAICS-3 industries* | | | |
| Within-industry adjustment | −3.3 | −0.0 | −5.3 |
| Between-industry reallocation | −0.7 | −0.4 | −1.0 |
| Residual | −0.9 | −0.6 | −1.0 |
| *NAICS-4 industries* | | | |
| Within-industry adjustment | −2.6 | +0.8 | −4.7 |
| Between-industry reallocation | −0.7 | −0.2 | −1.0 |
| Residual | −1.7 | −1.5 | −1.8 |

**Regional factors** As in the within-between industry decomposition of the labor share change, one could study regional factors. Such a regionally differential effect would be possible if firms sort into different regions according to their labor share. For example, states may provide tax incentives if firms open a new establishment in their state. If labor laws in these states are more lax, then workes may not be compensated as much as they are in other states. There is a vast array of reasons how some regions could have a different effect on the labor share than others. So we decompose the aggregate labor share decline into a within- and between region effect analogously to the within-between industry decomposition. Defining $j$ in equation (2) as one of the nine Census divisions, Table 4 displays the results.

Table 4: Within- vs. between-regional factors in the labor share decline

| Portions of labor share change | 1967-2007 | 1967-1982 | 1982-2007 |
|---|---|---|---|
| | (percentage point changes) | | |
| Aggregate labor share change | −4.9 | −0.9 | −7.3 |
| | | | |
| *Census divisions* | | | |
| Within-region adjustment | −4.1 | −0.1 | −6.5 |
| Between-region reallocation. | −0.3 | −0.6 | −0.1 |
| Residual | −0.6 | −0.2 | −0.8 |

As with the industry decomposition, most action occurs inside the regions rather than reflecting between-regional reallocation shifts: Of the 7.3 ppt decline since 1982, 6.5 ppt occur within divisions, whereas between-division reallocation accounts for a meek 0.1 ppt of the whole decline.

**Legal organization** As in the within-between industry and region decompositions of the labor share change, one could study the effect of the legal form of organization (LFO). Such an effect would be possible if firms with different legal form of organization lower their labor share. There has been lots of work on the impact of the 1986 tax reforms and how it gave rise to a new legal form of organization of businesses: the S-corporation. In fact, an increasing share of aggregate value added and employment is accounted for by that new type of firm. But this argument can be extended to any change of legal form of organization which comprise corporations, proprietorships, partnerships, co-operatives, etc. So we decompose the aggregate labor share decline into a within- and between legal-form-of-organizational effects analogously to the within-between industry and region decompositions. Defining $j$ in equation (2) as one of the legal form of organizations, Table 5 displays the results.

Table 5: Within- vs. between-legal form of organizational factors in the labor share decline

| Portions of labor share change | 1977-2012 | 1977-1982 | 1982-2012 |
|---|---|---|---|
| | (percentage point changes) | | |
| Aggregate labor share change | −5.7 | +2.3 | −6.2 |
| | | | |
| *Census divisions* | | | |
| Within-LFO adjustment | −6.3 | +1.1 | −6.6 |
| Between-LFO reallocation | +0.3 | −0.6 | +0.4 |
| Residual | +0.4 | +1.8 | +0.0 |

Unlike in other sectors – Zidar et al. find that the shift to S-corporations explains up to a third of the labor share decline in services – legal form of organization does not matter for the aggregate labor share decline in manufacturing.

## 2.5 The micro-level anatomy of the aggregate labor share decline

### 2.5.1 The labor share *increases* in most establishments

Next, we study the labor share dynamics for individual establishments in the U.S. manufacturing sector. To do so, we start by decomposing the aggregate labor share $\lambda_t$ as:

$$\lambda_t = \frac{\sum_i w_{it} L_{it}}{\sum_i Y_{it}} = \sum_i \lambda_{it} \omega_{it} \tag{3}$$

where $\lambda_{it}$ corresponds to the labor share of establishment $i$ at time $t$ and $\omega_{it} = Y_{it}/Y_t$ denotes the value added weight of establishment $i$. As a first step, we study the distribution of the micro-level labor shares $\lambda_{it}$ and the role played by reallocation through the weights $\omega_{it} = Y_{it}/Y_t$. Figure 4 plots several quantiles of the raw distribution of establishment-level labor shares $\lambda_{it}$ alongside the aggregate labor share in a given Census year. To do so, we construct the labor share for each

establishment and study the cross-sectional distribution in every Census year since 1967.

Figure 4: Aggregate and establishment-level labor shares

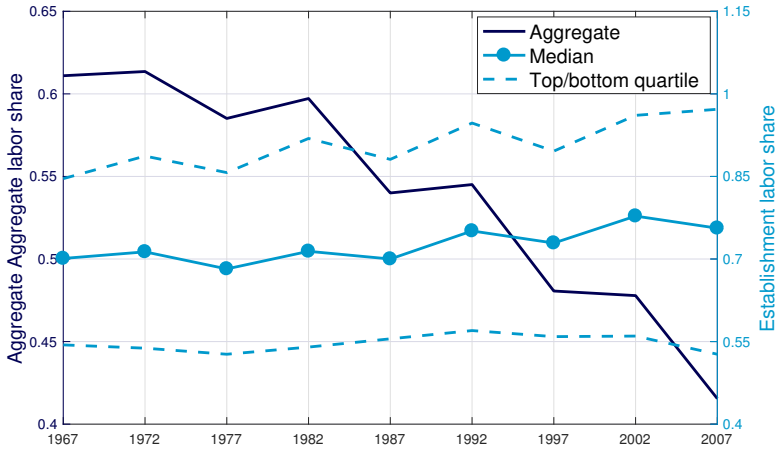


Figure 4 immediately reveals that the distribution of the establishment-level labor shares does not mimic the evolution of the aggregate labor share. While the aggregate labor share declines by 4.9 ppt per decade, the median labor share *increases* by 1.4 ppt per decade. The top and bottom quartiles strongly co-move with the median and increase as well. This evidence highlights diverging trends in the the labor shares at the aggregate and establishment level, a pattern particularly prevalent since 1982. An implication of this finding is that the aggregate labor share decline is not the result of a simple shift of the distribution of labor shares in individual establishments. Instead, our evidence points to the importance of reallocation as the main driver of the aggregate labor share dynamics. This is what we turn our attention to next.

### 2.5.2 The polarization of labor shares

The evolution of the unweighted distribution of plant-level labor shares reveals a second pattern: the distribution of (raw) labor shares becomes more polarized over time. Figure 4 suggests a widening of the raw labor share distribution around the median labor share. We examine the dynamics of the labor share distribution more carefully in Figure 5 by first plotting various dispersion statistics in the left panel. All of them are increasing over time: both the inter-quartile and inter-decile range are 40% more spread out in 2007 than they are in 1967, and the cross-sectional standard deviation rises strongly over time. In line with the quantile evidence from Figure 4, we notice that most of this widening of the distribution occurs at the right tail of the distribution, i.e. high-labor-share establishments pull farther away from the median relative to low-labor-share establishments.

To show the polarization of labor shares more directly, the right panel of Figure 5 plots their distributions at the beginning and end of our sample. Compared to the 1967 distribution of labor shares, the 2007 distribution has less weight in the middle and more mass both in the higher and the lower tails. This polarization can be also shown using the fourth normalized moment: excess

kurtosis falls from +0.24 in 1967 down to −0.62 in 2007. That the distribution becomes more platykurtic over time is typical of a polarization pattern.Of note is that fact that low-labor share establishments do not appear to take over the market, as there is no /emphprima facie evidence of strong exit dynamics. In fact, by 2007 there is a relatively larger portion of establishments with very high labor shares.

Figure 5: The polarization of labor shares

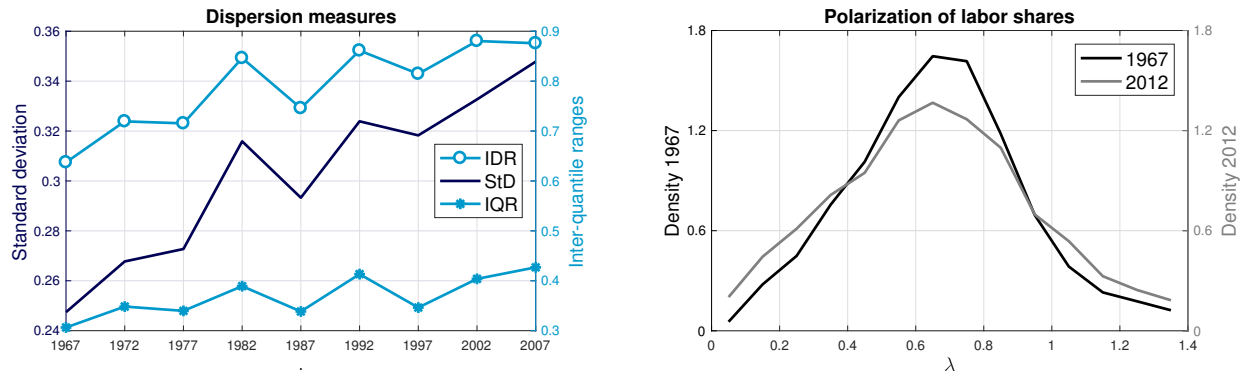


### 2.5.3 The importance of reallocation

So far, we have considered changes in the distribution of establishment-level labor shares, corresponding to the $\lambda_{it}$ terms in the decomposition of the aggregate labor share of equation (3). The upward trend of the raw labor share distribution, the increasing polarization and the widening of the distribution at the upper tail are in stark contrast to the decline of the aggregate labor share. Therefore, the $\omega_{it}$ terms in equation (3) must be the major force driving down the aggregate labor share, through a reallocation of value added to the lower tail of the labor share distribution. The more value added concentrates on low-labor share establishments, the lower the aggregate labor share holding the marginal distribution of labor shares constant. In the previous subsection, we saw that the latter is the case. We now consider where in the spectrum of labor shares value added is created. In other words, we divide the distribution of labor shares $\lambda$ into bins and consider the share of aggregate value added created in each bin and year. This will allow us to determine where in the labor share distribution most of the economic activity occurs, and how this pattern evolved over time.

#### The reallocation of output

Figure 6 displays both the distribution of labor shares $\lambda_{it}$ and value-added weights $\omega_{it}$ for each available Census year. The panels paint a stark picture: most of value added in 1967 is produced by establishments with a middle-of-the-road labor share between (50 and 80%); the value added weighted median is 62%. Over the following decades, however, the economic activity shifts grad-

Figure 6: Where in the labor share distribution does output get produced?

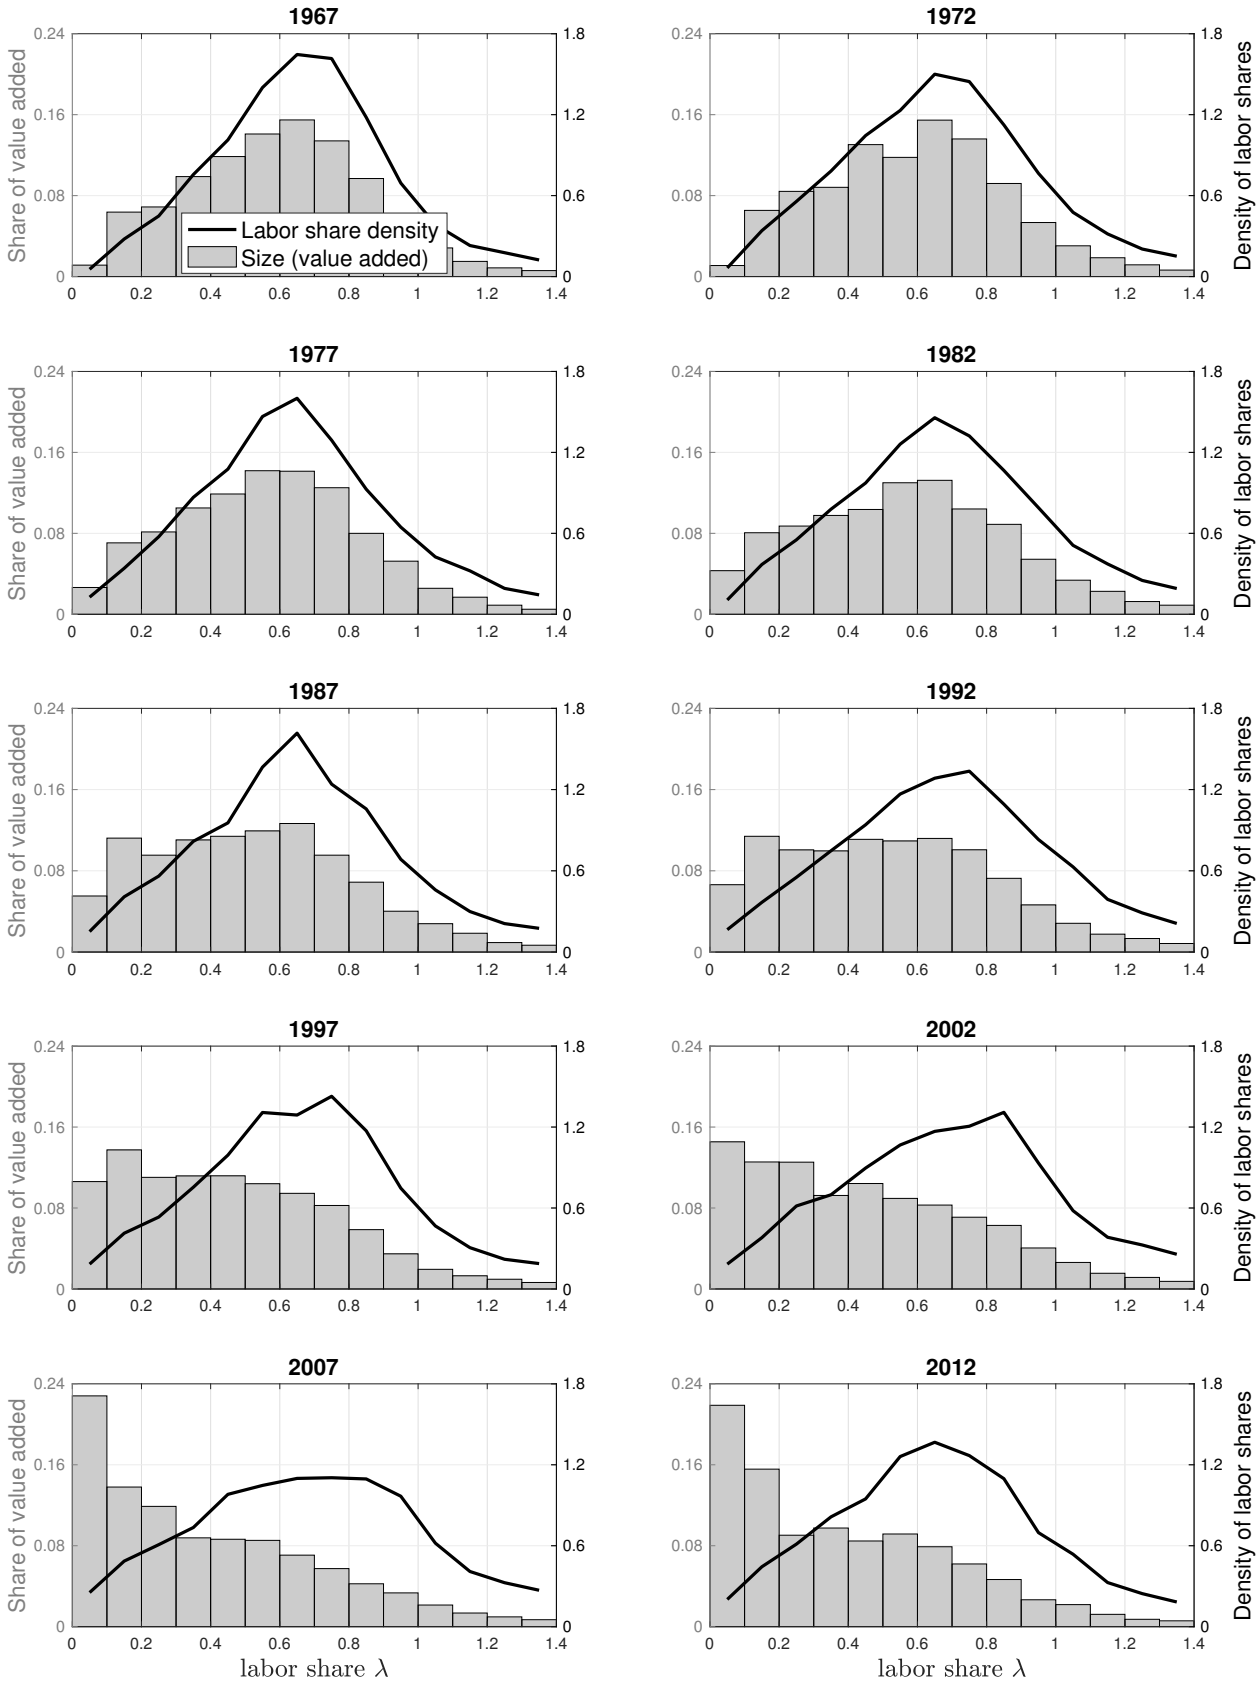

ually and persistently to the low-labor share spectrum. By 2007, half of aggregate value added if accounted for by establishments with a labor share less than 32%.

### Does labor get reallocated too?

The previous subsection showed how a massive reallocation of value added to establishments with a low labor share was the key driver of lowering the aggregate labor share. It is hence a natural question to ask if inputs get reallocated too. If establishments arrive at a low labor share due to low TFP, then it would be efficient to reallocate labor inputs to those firms. Knowing whether or not labor get similarly reallocated to low labor share as output does will help understand the drivers behind the reallocation of value added and the labor share decline.

Figure 7 displays the allocation of total employment across the spectrum of labor shares. One can see that in 1967 the bulk of workers are employed by firms with a medium labor share between 0.5 and 0.8. In the beginning of the sample, the allocation of labor and value added are similar in that both inputs and outputs concentrate on establishments with a medium labor share. Unlike value added which gets reallocated to the low end of the labor share distribution, however, the allocation of labor essentially doe not change. The bulk of workers still gets employed by establishments with a medium labor share. This means there is basically no reallocation of labor to the low end of the labor share distribution.

We executed the labor allocation exercise with three measures of employment which increasingly reflect the variable components of labor inputs: total employment, production workers and production hours. All three input measures yielded almost identical results with no reallocation of hours worked towards low labor share establishments. This is an important finding in light of Autor, Dorn, Katz, Patterson, and Reenen (2017b) who suggest a theory along the lines of Melitz and Ottaviano (2008) where firms pay a fixed labor cost (managers) and then hire variable labor (production workers) in order to produce. This model predicts that variable labor inputs monotonically increase in total factor productivity. This would mean that variable labor inputs such as production hours worked should be mostly accounted for by establishments at the low end of the labor share distribution. But Figure 7 shows that this is not the case. This means that interpreting their model with production hours as a variable labor input is inconsistent with the establishment-level evidence in manufacturing.

For completeness, we also report the distribution of relative wages and capital intensity in Figures 8 and 9. The former shows that low labor share establishments do pay a small wage premium – about 15% compared to high-labor share establishments within a given industry – which counteracts the labor share decline somewhat. In any case, the wage distribution within a given industry and year is remarkably compressed across the labor share spectrum and does not reflect the dynamics of labor productivity which rises a lot for low labor share establishments.

Figure 7: Where in the labor share distribution do workers get employed?

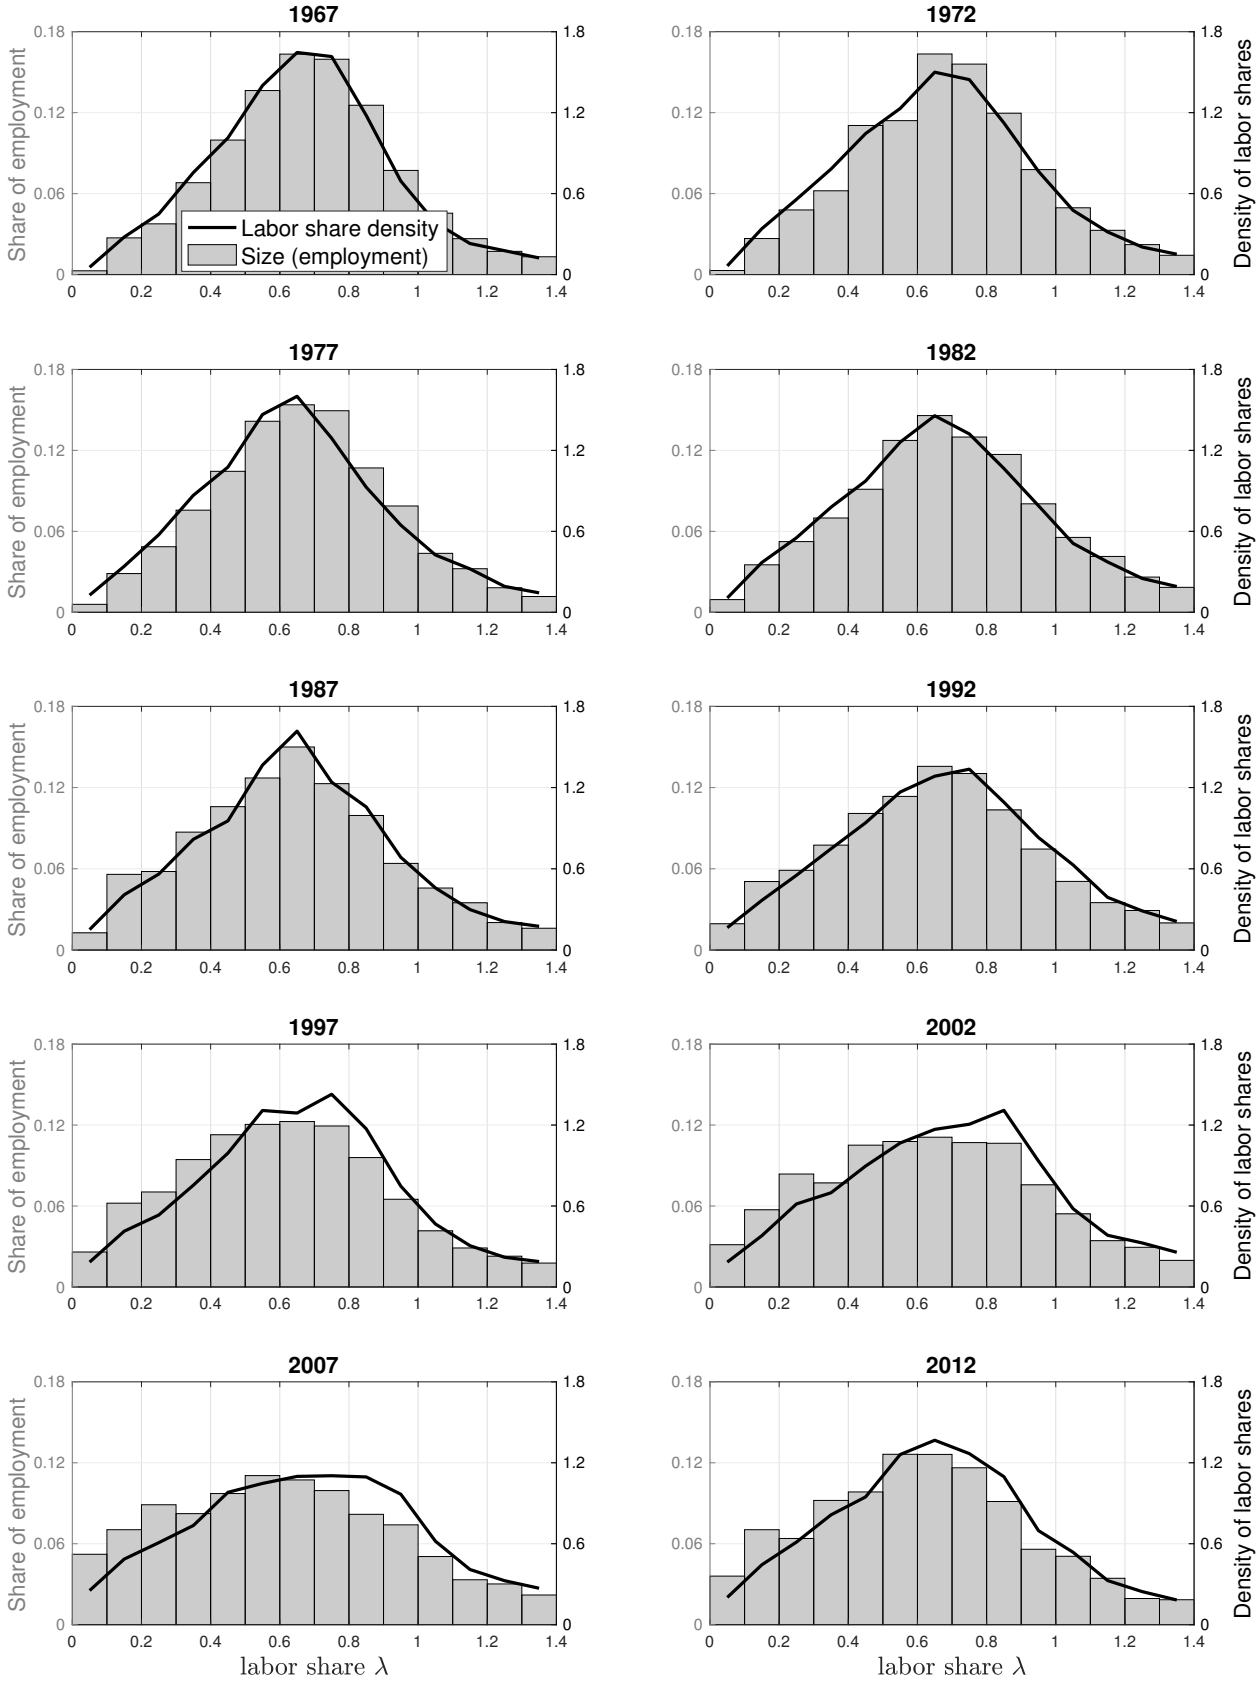

Figure 8: Labor share and wages

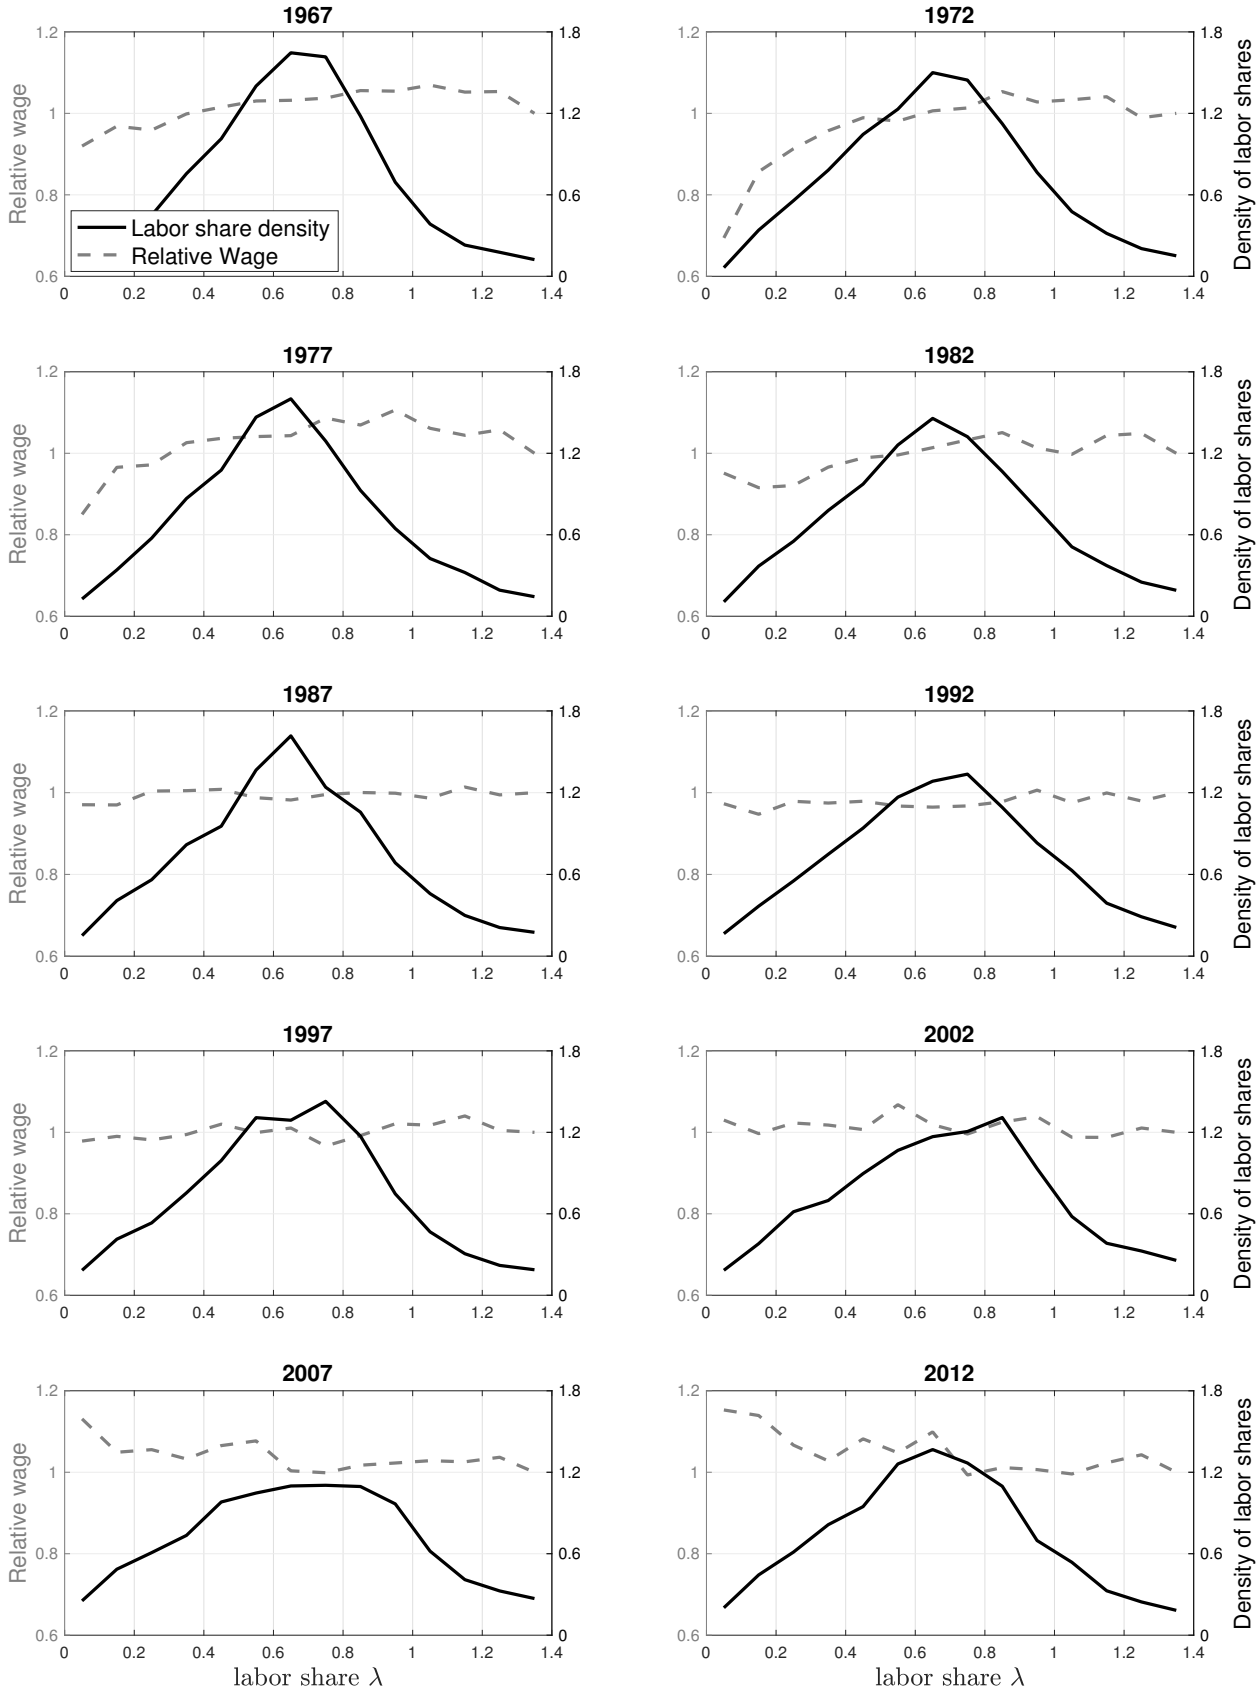

Figure 9: Labor share and capital intensity

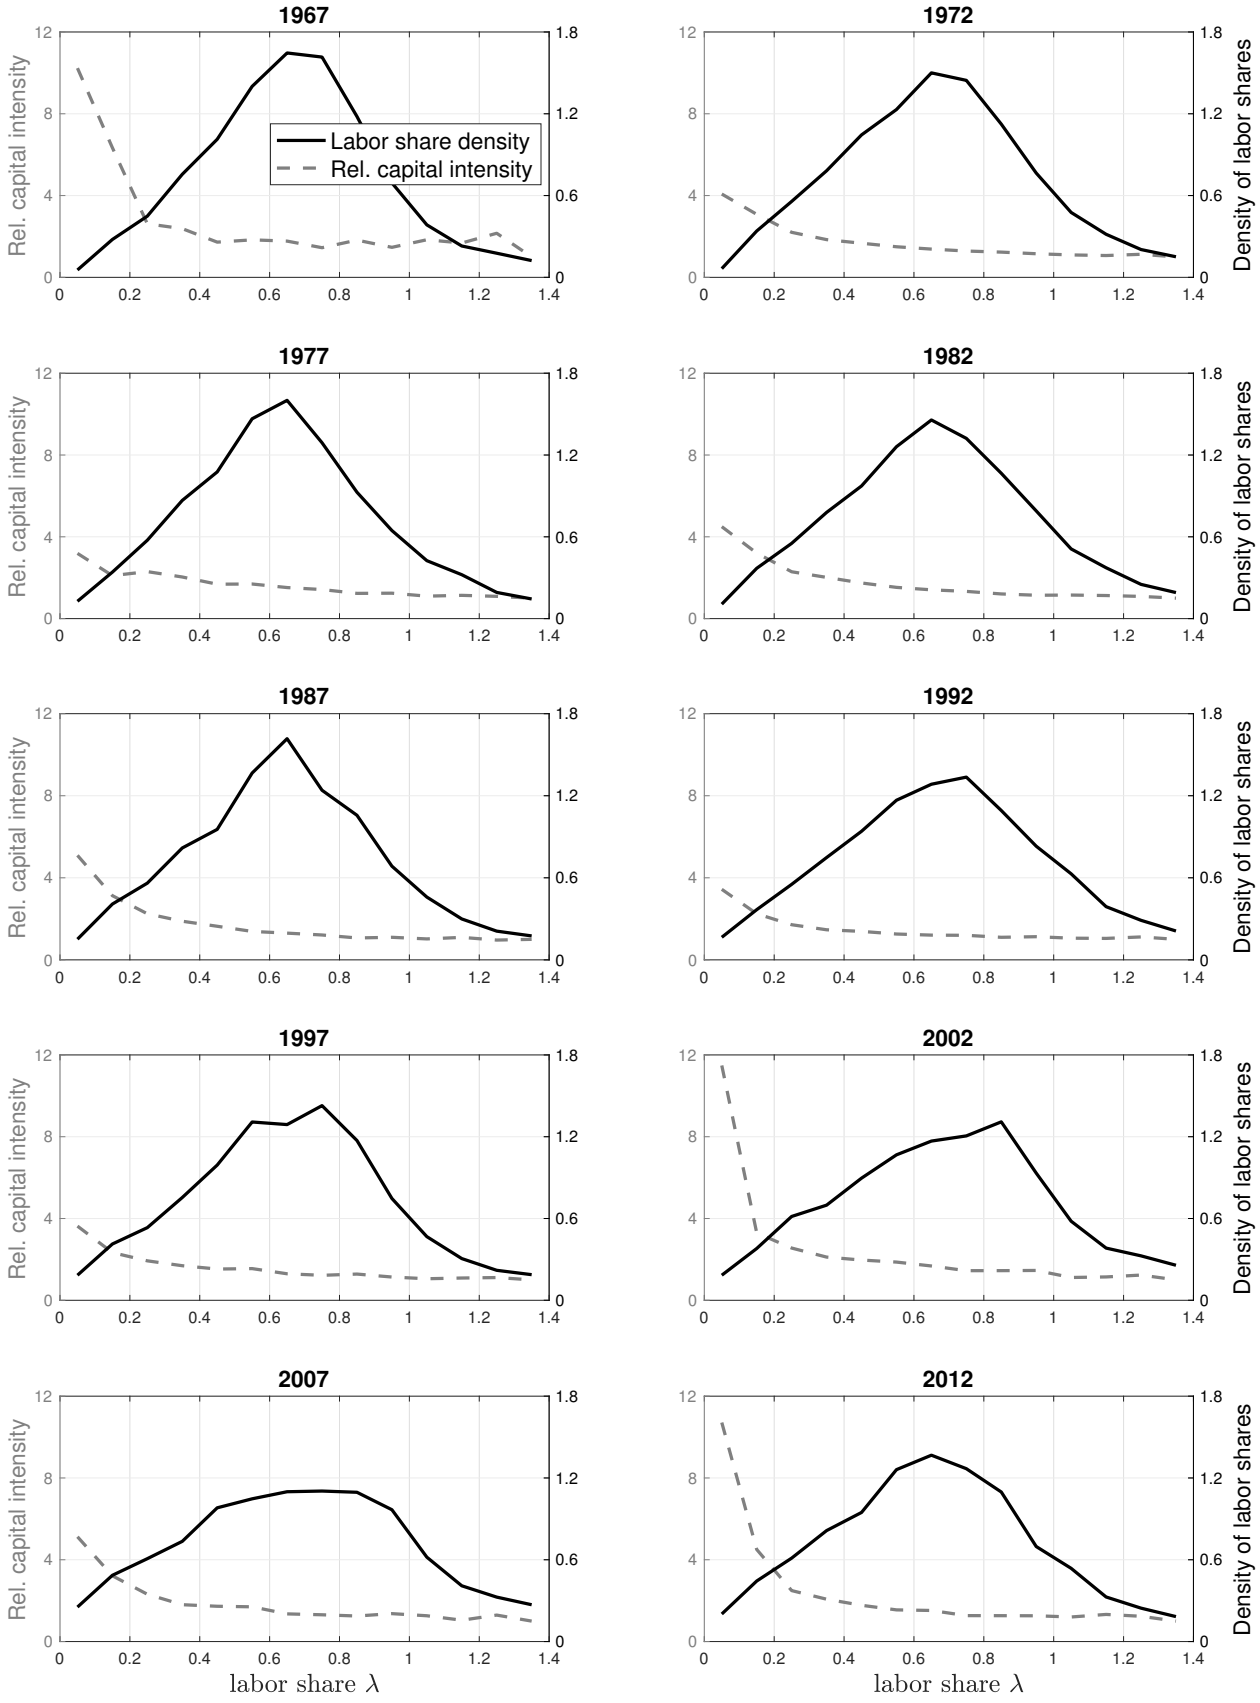

### 2.5.4 Reallocation along the extensive margin

There are several channels through which this concentration of economic activity across the labor share spectrum may occur: relatively high labor share establishments may exit; relatively low labor share establishments may enter; or there could be reallocation across continuing establishments. This section sheds light on the quantitative contributions of these three potential candidates.

First, we investigate whether the labor share of entrants and exiters is different than that of incumbents. To do so, whenever we observe a plant that enters, exits or changes owner, we compare its labor share to that of the other plants in the firm it just joined or left. The results are presented in Table 6, once we average these comparative statistics at the plant level. Generally speaking, we find that the differences in labor share between entering/exiting establishments and incumbent ones are sizable. Table 6 shows that exiting establishments are characterized by a labor share that is on average 5.6 ppt higher than that of their continuing peers. This difference is even stronger (13.2 ppt higher) when all establishments of the firm exit, while plants that are shut down by their continuing parent firm have a labor share that is 3.2 ppt higher. Establishments that are sold to a new owner, too, have a higher labor share of about 3.8 ppt. Maybe surprising is the fact that entering establishments have a labor share that is about 1 ppt higher than that of their incumbent peers. At first sight, this appears at odds with the notion that new entrants are high-productivity entrants. This seems at odds with the fact that the lower share of entering plants built by existing firms is considerably *lower* (about 4 ppt) than the average. Yet, there are potential reasons that may jointly explain these two facts. For one, it may reflect the fact that new establishments without an incumbent parent firm are too credit constrained to build capital, so they have to operate with an inefficiently low capital intensity. Alternatively, it could be the artefact of the important role played by learning.

Table 6: Labor shares by establishment type (percentage points)

| Labor share of ... | | | |
|---|---|---|---|
| new establishments in | | establishments acquired by | |
| new firms | exist. firms | new firms | exist. firms |
| *... relative to incumbent establishment operated by same firm* | | | |
| +7.3 | −4.0 | +5.2 | +0.8 |

| Labor share of ... | | | |
|---|---|---|---|
| exiting establishments in | | establishments sold by | |
| exiting firms | continuing firms | exiting firms | continuing firms |
| *... relative to continuing establishment operated by same firm* | | | |
| +13.2 | +3.2 | +3.9 | +3.7 |

Next, to assess the role played by each margin in the decline of the aggregate labor share, we

perform the following decomposition which is similar to the one in Melitz and Polanec (2015):

$$\begin{aligned}\Delta\lambda_t &= \sum_{i\in en}\lambda_{it}\omega_{it} + \sum_{i\in inc}\lambda_{it}\omega_{it} - \sum_{i\in ex}\lambda_{it-1}\omega_{it-1} - \sum_{i\in inc}\lambda_{it-1}\omega_{it-1} \\ &= \underbrace{\alpha_t^{en}(\lambda_t^{en}-\lambda_t^{inc})}_{\text{Contribution Entry}} \underbrace{-\alpha_{t-1}^{ex}(\lambda_{t-1}^{ex}-\lambda_{t-1}^{inc})}_{\text{Contribution Exit}} \underbrace{+\lambda_t^{inc}-\lambda_{t-1}^{inc}}_{\text{Change of incumbents}} \\ &= \underbrace{\alpha_t^{en}(\lambda_t^{en}-\lambda_t^{inc})}_{\text{Contribution Entry}} \underbrace{-\alpha_{t-1}^{ex}(\lambda_{t-1}^{ex}-\lambda_{t-1}^{inc})}_{\text{Contribution Exit}} \underbrace{+\Delta\overline{\lambda}_t^{inc}+\Delta Cov(\lambda_{it},\omega_{it})}_{\text{Change of incumbents}}\end{aligned}$$

where $\lambda_t^{en}$ denotes the aggregate labor share of the set of establishments that enter in period $t$, $\lambda_t^{inc}$ that of establishments that continued from the previous period and $\lambda_{t-1}^{ex}$ that of establishments which existed last period and then exited. $\alpha_t^{en}$ and $\alpha_{t-1}^{ex}$ are the weights of the entering and exiting populations which correspond to their respective shares of aggregate value added. Following Olley and Pakes (1996) and Melitz and Polanec (2015), the aggregate labor share of incumbents, $\lambda_t^{inc}$, can be expressed as the unweighted average labor share and the cross-establishment covariance between labor shares and value added.

Table 7: Contributions to the labor share decline: extensive vs. intensive margins

| Portions of labor share change | 1967-2007 | 1967-1982 | 1982-2007 |
|---|---|---|---|
| | (percentage point changes) | | |
| Aggregate labor share change | −4.9 | −0.9 | −7.3 |
| Contribution entry | +0.3 | −0.9 | +1.1 |
| ... into new firms | +1.2 | +0.3 | +1.7 |
| ... into incumbent firms | −0.9 | −1.3 | −0.7 |
| Contribution exit | −2.3 | −1.4 | −2.8 |
| ... from exiting firms | −1.6 | −0.8 | −2.1 |
| ... from continuing firms. | −0.7 | −0.5 | −0.7 |
| Contribution incumbents | −3.0 | +1.4 | −5.5 |
| ... from average change | **+4.2** | +6.8 | +2.6 |
| ... from covariance term | **−7.1** | −5.4 | −8.1 |

We display the contributions of each component of this exact decomposition in Table 7, over the entire sample as well as separately before and after then onset of the aggregate labor share decline in the early 1980s. Our results show that the impact of entrants on the labor share is rather small on average. Surprisingly, it actually contributes *positively* to the aggregate labor share in the second half of our sample. Even though new establishments built by incumbent firms lower the aggregate labor share in principle, this effect weakens in the second part of our sample. Both pieces of evidence clearly rule out entry as a leading candidate for the labor share decline.

Exit, in contrast, does impact the aggregate labor share negatively and more so since the 1980s. Exiting establishments are characterized by a labor share which is about 5 ppt higher than that of their continuing peers. They also account for enough of aggregate value added to matter quantitatively for the fall in the aggregate labor share: In the first part of the sample, exit contributes about 1.4 ppt a falling labor share; this jumps to 2.8 ppt of the aggregate labor share decline since 1982. Had this increased selection not occurred, the aggregate labor share decline would have only been 5.6 ppt instead of the observed 7.3 ppt. Interestingly, the increased impact of exit comes almost exclusively from firms shutting down entirely, while the closure of less productive plants by continuing firms contributes little to the overall decline.

The lion's share of the aggregate labor share dynamics, however, is accounted for by the intensive margin. This margin can be active if there are changes in the labor share in the population of incumbents, or through the reallocation of value added to relatively low labor share establishments away from their less productive peers. We find that this term is moderately positive with an upward contribution of 1.4 ppt per decade until 1982; this reflects the general right-ward shift of the labor share distribution documented above (more details on that below). But since the early 1980s, the intensive margin term becomes very negative (annual contribution of 5.5 ppt contribution per decade since 1982). In a nutshell, without this intensive margin, the decline in the aggregate labor share would be been only 2.2 instead of 7.3 ppt.

### 2.5.5 Reallocation along the intensive margin

The importance of the intensive margin warrants more investigation. As we mentioned earlier, it could be negative due to the reallocation of value added from relatively high- to relatively low-labor-share establishments, or because establishments have adjusted their labor share (or both). We modify a conventional shift-share decomposition[3] as follows:

$$\begin{aligned}\lambda_t^{inc} &= \int \lambda\omega_t(\lambda)dF(\lambda_t) \\ &= \int_{\lambda^{\min}}^{\lambda^1} \lambda\omega_t(\lambda)dF(\lambda_t) + ... + \int_{\lambda^{q-1}}^{\lambda^q} \lambda\omega_t(\lambda)dF(\lambda_t)... + \int_{\lambda^{99}}^{\lambda^{\max}} \lambda\omega_t(\lambda)dF(\lambda_t).\end{aligned}$$

[3] A traditional approach is the shift-share decomposition along plants (in the spirit of e.g. Baily et al. (1992)):

$$\Delta\lambda_t^{inc} = \underbrace{\sum_i \omega_{it-1}\Delta\lambda_{it}}_{\text{Shift: } +15.0\%} + \underbrace{\sum_i \Delta\omega_{it}\lambda_{it-1}}_{\text{Share: } +8.8\%} + \underbrace{\sum_i \Delta\omega_{it}\Delta\lambda_{it}}_{\text{Interaction: } -26.8\%} .$$

While intuitive, the difficulty arises from interpreting the interaction term should it be quantitatively relevant. It may be interpreted as "directed reallocation," i.e. reallocation to establishments that *lower* their labor share (as opposed to undirected reallocation (Share) which captures reallocation to establishments that have a relatively low labor share). Alternatively, the interaction term might be interpreted as a differential adjustment of establishments that are growing in size relative to their peers. In our application, the first two terms are strongly positive (!) and the last "interaction term" turns out to be enormously negative and dominating everything mechanically. That the interaction term has an opposite sign than the other two and dominates quantitatively makes this decomposition useless.

Then, the aggregate labor share change can be written as:

$$\Delta\lambda_{inc}^{t} = \int \lambda\omega_t(\lambda)dF(\lambda_t) - \int \lambda\omega_{t-1}(\lambda)dF(\lambda_{t-1})$$
$$\approx \underbrace{\sum_q \Delta\tilde{\lambda}_t^q \omega_{t-1}(\lambda_{t-1}^q)}_{\text{Adjustment (Shift)}} + \underbrace{\sum_q \tilde{\lambda}_{t-1}^q \Delta\omega_t(\lambda_t^q)}_{\text{Reallocation (Share)}} + \underbrace{\sum_q \Delta\tilde{\lambda}_t^q \Delta\omega_t(\lambda_t^q)}_{\text{Interaction term}}$$

where $q = 1, 2, ..., 100$ denotes a percentile in the raw labor share distribution. "Adjustment" refers to shifts of the whole labor share distribution where each percentile change is weighted by the share of aggregate value added, $\omega_t(\lambda)$. $\tilde{\lambda}_t^q$ refers to the average labor share in percentile $q$ and $\omega_t(\lambda^q) = \sum_{i:\lambda_{it}\in[\lambda_t^{q-1},\lambda_t^q)} \omega_{it}$. "Reallocation" refers to reallocation of market share across the labor share distribution (rather than plants). As always, the interaction term is harder to interpret; fortunately, it is almost zero in our application.

Figure 10 displays graphically the contribution of each term of the intensive margin, while Table 8 provides the detailed numbers. Two results stand out.

Figure 10: Contribution to the labor share decline

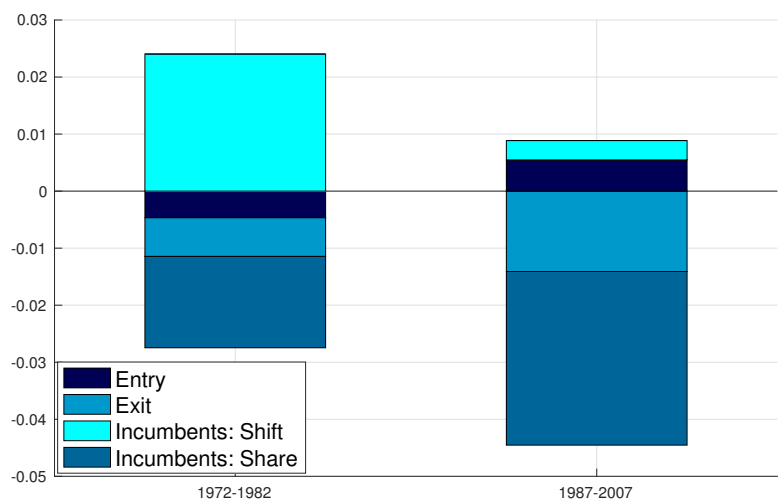


**Reallocation to low labor-share incumbents** The first main finding is that there is significant reallocation from high percentiles to low percentiles of the labor share distribution of incumbent establishments. We already knew from the decomposition in Section 2.5.3 that in the overall distribution, reallocation plays the key role rather than adjustment. Yet this need not be true for incumbents alone since the adjustment margin could be cancelled out through an opposite adjustment in the entry/exit margin. Our decomposition now allows us to make a stronger statement than the one following Figure 6. We find that reallocation always puts downward pressure of the labor share, but especially so after 1982: its contribution to the overall labor share decline almost

Table 8: Contributions to the labor share decline: components of the intensive margin

| Portions of labor share change | 1967-2007 | 1967-1982 | 1982-2007 |
|---|---|---|---|
| | (percentage point changes) | | |
| Aggregate labor share | −4.9 | −0.9 | −7.3 |
| | | | |
| Establishment entry | +0.3 | −0.9 | +1.1 |
| Establishment exit | −2.3 | −1.4 | −2.8 |
| | | | |
| Intensive margin | −3.0 | +1.4 | −5.5 |
| Adjustment (Shift) | +2.2 | +4.8 | +0.7 |
| Reallocation (Share) | −5.0 | −3.2 | −6.1 |
| Interaction term | −0.2 | −0.2 | −0.1 |

doubles in magnitude, from −3.2 ppt up to 1982 to −6.1 ppt afterwards.

**Reallocation to incumbents which decrease their labor share** Our second main result is that the adjustment margin is strongly positive before the 1980s and then drops to almost zero in the second part of the sample. This means that those portions of the raw labor share distribution that shifts downward (the left tail) accounts for an increasing share of economic activity. It is important to interpret this adjustment term jointly with the evidence of the raw labor share distribution. We documented above that the overall distribution is gradually shifting to the right over time, especially at the upper end of the distribution; the left tail of the distribution, on the other end, shifts slightly left. At the beginning of the sample, economic activity is evenly distributed across the upward-moving right tail and the downward-moving left tail of the distribution. This means that the adjustment term is positive as the entire distribution shifts right by an average 4.8 ppts per decade. Since the 1980s, however, more and more economic activity moves towards the low end of the labor share distribution, so that changes in this portion of the distribution no longer merely offset the upward movements in the right tail, but instead dominate them. The upshot is that for the purpose of its contribution to the decline in the aggregate labor share, the adjustment term makes it *look like* the distribution of plant-level labor shares shifted gradually and symmetrically to the left, while in fact it is the product of the interaction of both polarization and reallocation: While both the left tail declines and the right tail rises – thus leading to the polarization of labor shares –, any adjustment in the left tail is much more relevant because these establishments have grown more relative to those at the top tail of the labor share distribution. The small, but economically relevant shift at the left tail cancels the pervasive and large, but economically less relevant, rightward shift at the high end of the distribution: the adjustment term drop from 4.8 ppts down to 0.7 ppts per decade.

To summarize, we find that the change in the evolution of the aggregate labor share since the early 1980s can be decomposed in three main components: 1/6th of the downward pressure was

driven by more intense exit of relatively high-labor-share establishments; 1/3 by reallocation and 1/2 by the disproportionate growth of low-labor-share establishments. The contribution from the entry of relatively high labor share plants is *positive* but small.

In the next section, we turn to the analysis of these incumbent establishments that contributed most to the decline in the aggregate manufacturing labor share.

# 3 The dynamics of plant-level labor share

So far, we have shown using plant-level data that the decline in the aggregate manufacturing labor share had not been driven by a fall in the distribution of labor shares across establishments over time, but instead a dramatic reallocation of value added towards the lower end of this distribution. At first sight, this may seem to indicate that a small number of very productive "superstar" plants slowly but surely took over the manufacturing sector, leveraging their size to gain ever larger market shares. In this section, we show that the story is much more subtle. Our analysis of the dynamics of the labor share at the *micro* level exposes an environment that looks more like a "tilted lottery" than a "superstar economy."

## 3.1 The transient nature of plant-level labor shares

Our objective in this section is to investigate the persistence of the labor share at the plant level. We first document the serial correlation of labor share over our sample period before focusing our attention on dynamics for the establishments with the lowest labor shares.

### 3.1.1 Micro-level serial correlation of the labor share

As a first, simple exercise, we estimate a basic autoregressive process:

$$\lambda_{it} = \alpha + \rho\lambda_{it-k} + u_i + \varepsilon_{it} \tag{4}$$

where $k = 5$ for the quinquennial Census panel and $k = 1$ for the annual ASM panel.

The estimated coefficients are reported in Table 9. Running a panel regression with plant fixed effects, we find that the estimated $\rho$ is not significantly different than zero in the weighted regression using data from the Census of Manufactures (column (I)). This very low degree of persistence at five-year intervals is confirmed by the results from the Annual Survey of Manufactures: on a yearly basis, the autoregressive coefficient using value-added weights is also quite low, at 0.466 (column (III)). This implies that out of a 10ppt innovation in the labor share, only 4.6ppt remain after a year and 0.2ppt after five years, which is similar to the 0.1ppt from the Census estimate (column (I)).

It is worth noting that the unweighted results (columns (II) and (IV)) show less persistent labor shares: large establishments tend to have wage bills that move more in lockstep with value added, presumably because innovations to either are lower. In fact, the unweighted regression in

Table 9: Volatility of plant-level labor shares

| | (I) | (II) | (III) | (IV) |
|---|---|---|---|---|
| $\alpha$ | 0.485*** | 0.629*** | 0.389*** | 0.381*** |
| | (0.035) | (0.068) | (0.069) | (0.070) |
| $\rho$ | 0.011 | -0.110*** | 0.466*** | 0.328*** |
| | (0.012) | (0.007) | (0.011) | (0.010) |
| $N$ | 420k | 420k | 1,222k | 1,222k |
| Sample | CMF | CMF | ASM | ASM |
| VA weights | yes | no | yes | no |

*Note:* The table reports the results from a panel regression (4) in the Census of Manufactures (columns (I) and (II)) and the Annual Survey of Manufactures (columns (I) and (II)). The former sample focuses on plants that are observed at least four times out of the ten Census waves in order to obtain reliable autoregressive estimates. Weights are a plant's average market share within a 3-digit NAICS industry throughout its life time. Standard errors are clustered at the 3-digit NAICS industry level.

the Census data estimates a slightly negative autocorrelation coefficient. All in all, the unweighted regressions confirm that plant-level labor shares exhibit very little persistence.

### 3.1.2 The labor share dynamics of "hyper productive plants"

Next, we focus our attention on what we term "hyper productive plants" ($HP$ plants), establishments that are in the lowest quintile of the labor share distribution in a given 3-digit NAICS industry and year. As we have shown earlier, these hyper productive establishments have played a crucial role in the fall of the aggregate labor share over the 1967-2012 period, while accounting for a dramatically increasing portion of manufacturing value added.

There are two ways $HP$ plants can start to dominate the economy: they may have always had a low labor share and just grow their market share or they may have a temporarily low labor and high market share. The former would be consistent with a "superstar economy," the latter with a "lottery economy." These two polar cases have vastly different implications for our understanding of the aggregate labor share decline. To discriminate between them, we start by asking a simple question: conditional on a plant's labor share at $t-5$, what is the probability that it has $HP$ status at time $t$? If plant-level labor share is highly persistent like in a superstar economy, this probability should be equal to 100% for plants that were $HP$ at time $t-5$. On the polar opposite, in the lottery economy, this number would be 20% (one fifth as we consider quintiles) if plant-level labor shares were to be completely transient within an industry.

Table 10 shows that over our sample period, the probability for a plant to retain $HP$ status from Census year to Census year (a 5-year window) is only 41%. While this is higher than under the assumption that $HP$ status is perfectly random (20%), the transition probability confirms that labor share at the plant level is surprisingly transient, even for the most productive plants. This is consistent with the generally low persistence of labor shares from our auto-regressive estimates (see Table 9).

Table 10: The odds of becoming $HP$ by previous $HP$ status

| *Panel A. $HP$ is raw lowest quintile of $\lambda$* | | |
|---|---|---|
| | Non-$HP_t$ | $HP_t$ |
| Non-$HP_{t-5}$ | 0.854 | 0.146 |
| $HP_{t-5}$ | 0.585 | 0.415 |
| *Panel B. $\widetilde{HP}$ is lowest $\omega$-weighted quintile of $\lambda$* | | |
| | Non-$\widetilde{HP}_t$ | $\widetilde{HP}_t$ |
| Non-$\widetilde{HP}_{t-5}$ | 0.949 | 0.051 |
| $\widetilde{HP}_{t-5}$ | 0.719 | 0.281 |

The evidence in Panel A. of Table 10 was obtained by counting each plant as one unit, irrespective of its size. For Panel B., the $\widetilde{HP}$ plants are those with the lowest labor shares *and* whose cumulative value added is equal to one fifth of the total value added within a specific industry and year. We find that only 28.1% of the value added accounted for by $\widetilde{HP}_{t-5}$ remains $\widetilde{HP}$ while the remaining three quarters turn over to the $Non-\widetilde{HP}$ population. This means that among the very large, very low labor share plants there is even more turnover than among the $HP$ population. In sum, both Panel A. and Panel B. indicate that being a hyper-productive plant is a rather temporary phenomenon.

In order to confirm this finding, next we estimate and plot the typical time-series pattern of a time-$t$ $HP$ plant in the years before and after period $t$. To operationalize the time series pattern of HP labor shares, we first run the following regressions:

$$\text{for Census: } \lambda_{it+k} = c_k + \beta_k^{HP}\mathbb{I}\{i \in HP_t\} + \gamma_k X_{it} + \varepsilon_{it} \qquad k \in \{-5, 0, 5\}$$
$$\text{for ASM: } \lambda_{it+k} = c_k + \beta_k^{HP}\mathbb{I}\{i \in HP_t\} + \gamma_k X_{it} + \varepsilon_{it} \qquad k \in \{-5, -4, ..., 4, 5\}$$

where $X_{it}$ is a set of industry and regional controls. The $\beta_k^{HP}$ coefficients capture how much lower the average labor share of time-$t$ $HP$ plants is compared to non-$HP$ plants $k$ years before or after. For example, $\beta_0^{HP}$ estimates the contemporaneous labor share gap between $HP$ and non-$HP$ plants, while $\beta_{-5}^{HP}$ indicates how much lower (if at all) the labor share of $HP$ plants is compared to their peers five years before these plants attain $HP$ status.

The left panel of Figure 11 summarizes visually these estimates. It shows a very striking V-shaped pattern: while time-$t$ $HP$ plants have, by definition, a much lower average labor share than their peers in period $t$ – almost 40 ppts –, this difference is much less pronounced for the periods before and after. In other words, plants in the bottom quintile of the labor share distribution today were not in the past and will not be in the future significantly different than other establishments.

Figure 11: The temporary fall and rise of HP plants' labor shares

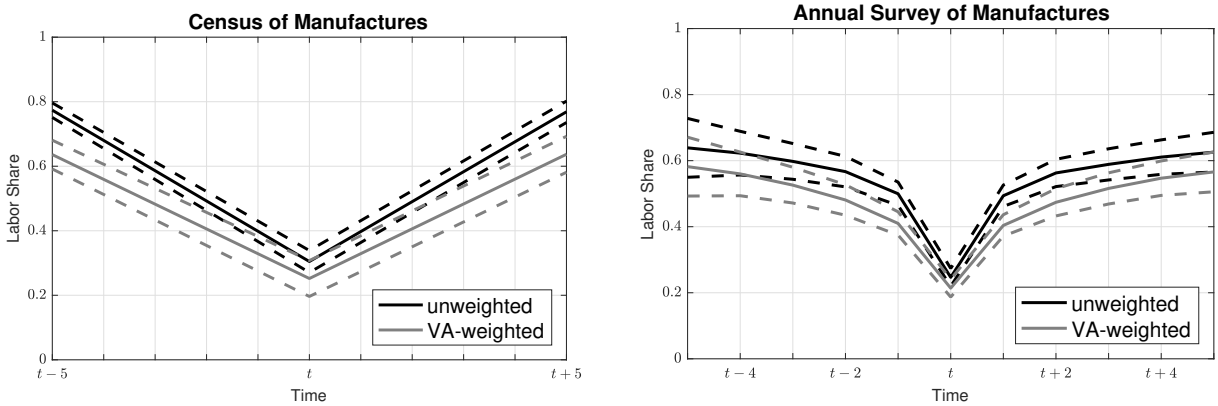


*Note:* Evolution of the labor share of the average $HP$ plant before and after the year it is in $HP$ status. Left panel: Census data, right panel: ASM data.

**$HP$ status and measurement error.** One potential concern is that the low persistence of the labor share is an artifact of measurement error. Under this scenario, $HP_t$ plants would simply be establishments that experienced important mismeasurement at time $t$, yet whose fundamentals were not any different than the typical plant in the population. This would mechanically give rise to the sharp V-shapes of Figure 12. To alleviate this concern, we turn our attention to the Annual Survey of Manufactures (ASM) sample. While this yearly dataset does not capture the population of manufacturing plants, its aggregate labor share dynamics are very similar to those of the Census. Crucially, its yearly frequency allows us to more easily disentangle signal from noise: if $HP$ status were merely driven by idiosyncratic measurement error, we would expect plants that are $HP$ plants to look on average like their non-$HP$ peers not only five years before and after (Census frequency), but also in the years directly following and preceding year $t$ (ASM frequency). The right panel of Figure 11 confirms once again the transient nature of the labor share. However, while the trough at $t$ is unmistakable, notice that the drop in the labor share starts occurring in the preceding years, and does not fully recover in the immediate subsequent years. More specifically, the labor share of current $HP$ plants is about 15ppts below that of their peers at $t-1$ and $t+1$. This confirms our earlier finding in Table 9 that while the persistence of the labor share was very low at 5-year intervals, it was significantly higher at the 1-year frequency. All in all, our evidence appears to confirm that $HP$ status is not merely an artifact of measurement error.

### 3.1.3 The evolution of the V-shaped pattern over time

The results so far were obtained by pooling together all years. Next, we investigate how plant-level labor share dynamics have evolved over time. Figure 12 illustrates the evolution of the depth of the V-shape pattern, but this time estimated for each separate Census year. Two things are worth noticing. First, the transient nature of the plant-level labor share is not a new phenomenon and is

present in every year of our sample. Second, the depth of the V-shape appears to have increased over time: for example, while $HP$ plants (VA weighted) had a labor share 35 percentage points lower than their peers in 1967, the gap had grown to 50 percentage points by 2007.

Figure 12: Labor share of $HP$ plants relative to peers

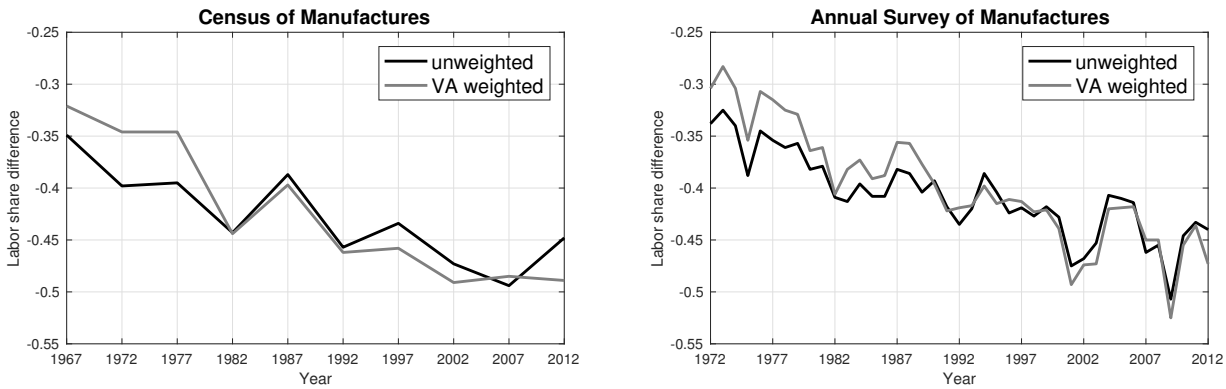


## 3.2 The importance of "hyper productive plants": Size and selection into HP status

In the previous section, we found not only that $HP$ status is highly transient, but also that $HP$ plants have become increasingly different than their peers. Yet, the deepening of the V-shape pattern that we documented cannot per se explain the significant decline in the aggregate labor share. After all, in every period a similar number of plants are in the $HP$ and non-$HP$ portions of labor share distribution. If the non-$HP$ portion experiences an inverse V-shaped pattern, then their respective contributions are more or less cancelling each other. Since the aggregate labor share depends not only on the plant-level labor share distribution, but also the joint labor share-size distribution, the *type* of plants that are selected to become $HP$ matters. In this section, we show that selection into $HP$ status has become increasingly correlated with size.

As a first exercise, we run the following regression for each Census year:

$$\mathbb{I}\left\{HP_{it+5}\right\} = \beta_1 \mathbb{I}\left\{\omega_{it} \in [0, \omega_t^{q1}]\right\} + ... + \beta_5 \mathbb{I}\left\{\omega_{it} \in [\omega_t^{q4}, \omega_t^{q5}]\right\} + X_{it} + \varepsilon_{it} \tag{5}$$

where $\omega_t^{q1}, ..., \omega_t^{q5}$ are the five market-share quintiles within an industry in year $t$. $X_{it}$ is a vector of industry and regional dummies. Each $\beta$ coefficient captures the probability that a time-$t+1$ $HP$ plant came from a specific size quintiles at time $t$. If the size of a plant today has no impact on its likelihood to be in the bottom labor share quintile tomorrow, then the coefficients should all be equal to 0.2. vcvd

Table 11 indicates that while market share is not a particularly strong predictor of future labor share in the early part of the sample, size becomes increasingly correlated with future $HP$ status with time. For example, of the plants that were in the lowest labor-share quintile in 1972, 26% came

Table 11: The odds of becoming future $HP$ by current size quintile

| | VA-weighted Quintile | | | | |
|---|---|---|---|---|---|
| Year | 1 | 2 | 3 | 4 | 5 |
| 1967 | 0.14 | 0.16 | 0.21 | 0.25 | 0.25 |
| 1972 | 0.12 | 0.17 | 0.20 | 0.24 | 0.26 |
| 1977 | 0.12 | 0.16 | 0.21 | 0.24 | 0.27 |
| 1982 | 0.09 | 0.15 | 0.21 | 0.27 | 0.29 |
| 1987 | 0.09 | 0.14 | 0.19 | 0.25 | 0.32 |
| 1992 | 0.08 | 0.14 | 0.21 | 0.27 | 0.31 |
| 1997 | 0.09 | 0.14 | 0.20 | 0.25 | 0.31 |
| 2002 | 0.09 | 0.14 | 0.19 | 0.26 | 0.33 |
| 2007 | 0.08 | 0.13 | 0.19 | 0.26 | 0.34 |

| | Unweighted Quintile | | | | |
|---|---|---|---|---|---|
| Year | 1 | 2 | 3 | 4 | 5 |
| 1967 | 0.20 | 0.18 | 0.17 | 0.18 | 0.26 |
| 1972 | 0.18 | 0.16 | 0.17 | 0.21 | 0.29 |
| 1977 | 0.20 | 0.17 | 0.16 | 0.19 | 0.28 |
| 1982 | 0.14 | 0.15 | 0.16 | 0.21 | 0.34 |
| 1987 | 0.15 | 0.16 | 0.17 | 0.19 | 0.32 |
| 1992 | 0.13 | 0.14 | 0.17 | 0.21 | 0.35 |
| 1997 | 0.14 | 0.15 | 0.18 | 0.21 | 0.33 |
| 2002 | 0.13 | 0.15 | 0.17 | 0.20 | 0.34 |
| 2007 | 0.13 | 0.14 | 0.17 | 0.20 | 0.36 |

*Note:* Regression results from Equation (5), run separately by year=1967, ..., 2007. Size quintiles are defined along plant $i$'s market share in year $t$ within its 3-digit NAICS industry $j$ ($\omega_{ijt}/\sum_{i\in j}\omega_{ijt}$) and range from smallest (Quintile 1) to largest (Quintile 5).

from the largest 1967 size quintile compared to 18% from the second lowest quintile (VA-weighted numbers). For the 2012 $HP$ plants however, the same proportions were 36% and 14%, respectively, five years earlier. The pattern is similar if we do not weigh plants by their value added: the share of $HP$ plants that were from the first size quartile five years ago drops from 20% to 13% over the sample period, while the proportion from the largest size quintile rises from 26% to 36%.

### 3.3 Counterfactuals

— TBD —

## 4 The changing labor response to technology shocks

The perhaps most surprising fact about "hyper-productive plants" is their non-response on the labor side to profitability shocks. Standard models would predict that high-productivity establishments expand their workforce or at least increase their wages. The previous analysis showed that none of that was the case. This conforms well with work by Decker, Haltiwanger, Jarmin, and Miranda (2017a,b); Cooper, Haltiwanger, and Willis (2017) who have shown that the responsiveness of shocks has declined. Related work by Ilut, Kehrig, and Schneider (forthcoming) shows that establishment hiring responses to negative shocks are about 1.5 times as strong as hiring responses to positive shocks to future profitability. These results combined can provide an explanation for the declining labor share. Taking the hiring asymmetry from Ilut, Kehrig, and Schneider (forthcoming) this may be the case if more and more establishments find themselves in the right tail of the labor share distribution and/or if the sensitivity of hiring to profitability shocks declined in general.

To find out about these conjectures, we study how hiring, the numerator of the labor share, responds to profitability shocks that affect the denominator. We follow the empirical setup in Ilut,

Kehrig, and Schneider (forthcoming) and non-parametrically estimate net hiring as a function of TFP shocks

$$\hat{n}_{it} = f(z_{it}) + \varepsilon_{it}$$

where $z_{it}$ is the TFP shock of establishment $i$ in period $t$ and $n_{it}$ is its employment growth. We estimate this hiring function by 3-digit NAICS industry, size and decade. The former two characteristics control for industry-specific and size-dependent differences in hiring, and we aggregate the estimated hiring responses to obtain the typical hiring response of heterogeneous establishments within an industry and decade. We display our non-parametric estimates in Figure 13 which allows us to see how the relationship between hiring and TFP shocks changes. Figure 13 shows that hiring

Figure 13: The increasingly asymmetric productivity-hiring relationship

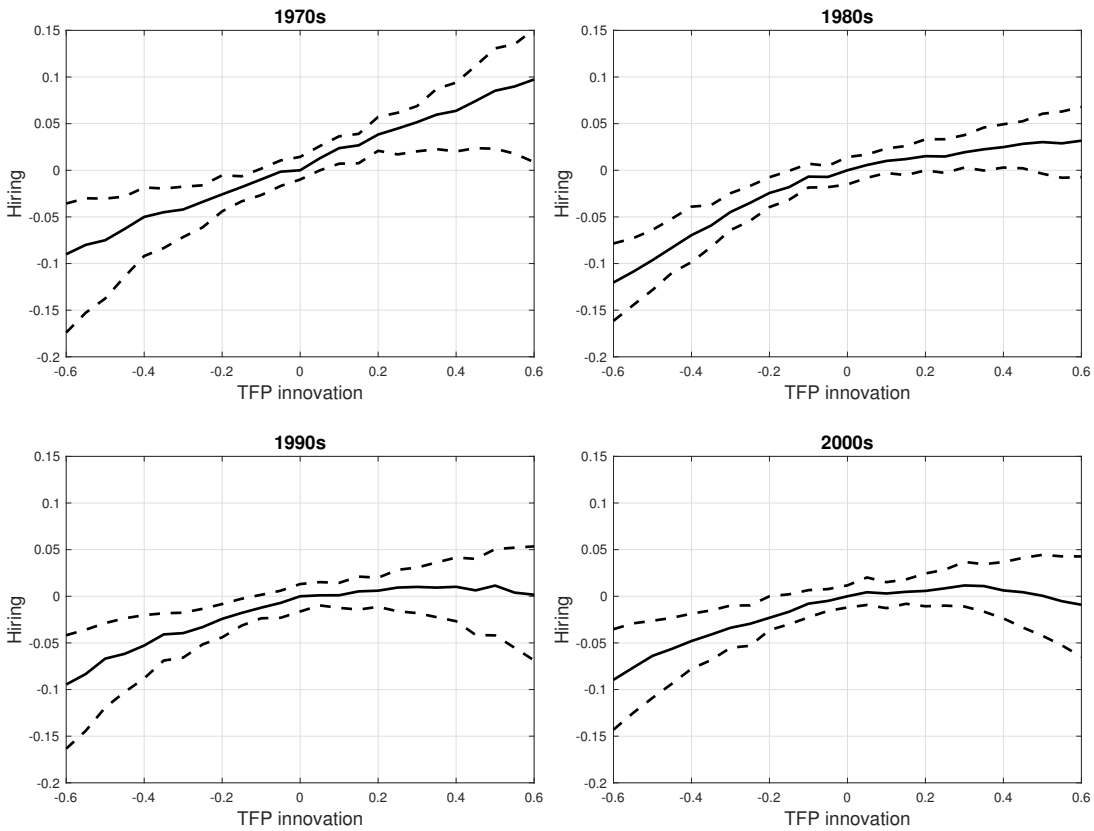


becomes increasingly asymmetric over time. While hiring in the 1970s is almost linear in technology shocks, it becomes asymmetric in the 1980s. This is because hiring after positive TFP shocks is very weak, almost insignificant at the 5% level at a two-standard deviation TFP shock ($z \approx +0.4$). Negative TFP shocks, in contrast, lead to significant firing as in the 1970s. The bottom two panels of Figure 13 show that the hiring-TFP shock relationship becomes even more asymmetric in the 1990s and 2000s.

At the same time, the TFP shock distribution changes. As Kehrig (2011) has shown, the long-

run cross-sectional dispersion of TFP levels becomes wider with every recession since the 1980/82 recession. This means there are more firms in the tails of the TFP innovation distribution in Figure 13. This increasing polarization of TFP shocks means that the asymmetry bears out more strongly towards the 2000s than in the 1970s.

Why don't high-TFP forms not hire anymore in the 2000s relative to their hiring response in the 1970s? The dynamics of the labor share in the previous section give an answer why. While the distribution of establishment TFP gets wider oder time, having high TFP and thus a low labor share in one period also means it is likely to have a low TFP again next period. If hiring is a dynamic choice in that it is hard to fire workers or costly to train them in the first place, the establishment might be reluctant to act on a very positive realization of TFP.

# 5 Conclusion

A large literature has recently documented and studied the decline in the labor share, both at the national and sectoral levels. In this paper, we dissect the underlying dynamics behind this phenomenon by using plant-level data for the U.S. manufacturing sector between 1967 and 2007. We first document a startling fact: while the aggregate labor share declined by almost 5 ppts per decade starting in the early 1980s, the labor share of the median plant *rose* over the same time period.

This apparent disconnect is due to two main factors: a drastic reallocation of production from high labor share plants towards their low labor share peers, as well as an additional downward adjustment of the labor share of these latter establishements, which we label *hyperproductive plants*. We show that these plants were able to decrease their labor share by increasing their value added drastically, without raising wages or employment significantly more than other plants. In ongoing work, we investigate further the characteristics of these hyperproductive plants in order to better understand the channels that are behind the decline in the aggregate labor share.

==> Why it matters Why look? If it's permanent, you'd expect that HP plants hire/invest, get bigger and hteir LS comes down eventually. If it's temporary, they will not invest/hire because they anticipate their LS will come down again, but completely different reasons that matter for growth. Also, very different implications regarding the impact of concentration on competition.

## References


Daron Acemoglu and David Autor. Skills, tasks and technologies: Implications for employment and earnings. In David Card and Orley Ashenfelter, editors, *Handbook of Labor Economics*, volume 4, Part B, chapter 12, pages 1043–1171. Elsevier, Amsterdam, 2011.

Daron Acemoglu and Pascual Restrepo. The race between machine and man: Implications of technology for growth, factor shares and employment. *NBER Working Paper No. 22252*, 2016.

Daron Acemoglu and Pascual Restrepo. Robots and jobs: Evidence from US labor markets. *NBER Working Paper No. 23285*, 2017.

Francisco Alvarez-Cuadrado, Ngo Van Long, and Markus Poschke. Capital-labor substitution, structural change and the labor income share. *IZA Discussion Paper No. 8941*, March 2015.

David Autor, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen. Concentrating on the fall of the labor share. *American Economic Review Papers and Proceedings*, 107 (5), May 2017a.

David Autor, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen. The fall of the labor share and the rise of superstar firms. *NBER Working Paper No. 23396*, 2017b.

Martin N. Baily, Charles Hulten, and David Campbell. Productivity dynamics in manufacturing plants. *Brookings Papers on Economic Activity*, 1992:187–267, 1992.

David Rezza Baqaee and Emmanuel Farhi. Productivity and misallocation in general equilibrium. *Working Paper*, 2017.

Simcha Barkai. Declining labor and capital shares. *Working Paper*, 2017.

Daniel Berkowitz, Hong May, and Shuichiro Nishioka. Does capital-labor substitution or do institutions explain declining labor shares? *Working Paper*, 2017.

Petri Böckerman and Mika Maliranta. Globalization, creative destruction, and labour share change: Evidence on the determinants and mechanisms from longitudinal plant-level data. *Oxford Economic Papers*, 64(2):259–280, April 2012.

Christoph E. Boehm, Aaron Flaaen, and Nitya Pandalai-Nayar. Multinationals, offshoring, and the decline of U.S. manufacturing. *Working Paper*, 2015.

Benjamin Bridgman. Is labor's loss capital's gain? gross versus net labor shares. *Working Paper*, 2014.

Erik Brynjolfsson, Andrew McAfee, Michael Sorell, and Feng Zhu. Scale without mass: Business process replication and industry dynamics. *Harvard Business School Technology & Operations Mgt. Unit Research Paper No. 07-016*, September 2008.

Russell W. Cooper, John C. Haltiwanger, and Jonathan L. Willis. Declining dynamism at the establishment level. *SED Meeting Paper*, 2017.

Carol A. Corrado, Charles Hulten, and Daniel Sichel. Measuring capital and technology: An expanded framework. In Carol Corrado, John Haltiwanger, and Daniel Sichel, editors, *Measuring Capital in the New Economy*, volume 65 of *Studies in Income and Wealth*, chapter 1, pages 11–46. University of Chicago Press, 2005.

Carol A. Corrado, Charles Hulten, and Daniel Sichel. Intangible capital and US economic growth. *Review of Income and Wealth*, 55(3):661–685, September 2009.

Steven J. Davis, John C. Haltiwanger, Ron Jarmin, and Javier Miranda. Volatility and dispersion in business growth rates: Publicly traded versus privately held firms. In Daron Acemoglu, Kenneth Rogoff, and Michael Woodford, editors, *NBER Macroeconomics Annual*, volume 21, chapter 2, pages 107–179. University of Chicago Press, May 2006.

Ryan A. Decker, John Haltiwanger, Ron S. Jarmin, and Javier Miranda. Declining dynamism, allocative effciency, and the productivity slowdown. *American Economic Review Papers and Proceedings*, 107(5):322–326, May 2017a.

Ryan A. Decker, John Haltiwanger, Ron S. Jarmin, and Javier Miranda. Changing business dynamism and productivity: Shocks vs. responsiveness. *Working Paper*, 2017b.

Maya Eden and Paul Gaggl. On the welfare implications of automation. *Review of Economic Dynamics*, 29:15–43, July 2018.

Jan Eeckhout and Jan De Loecker. The rise of market power and the macroeconomic implications. *NBER Working Paper No. 23687*, 2017.

Michael W. L. Elsby, Bart Hobijn, and Ayşegül Şahin. The decline of the U.S. labor share. *Brookings Papers on Economic Activity*, pages 1–63, Fall 2013.

Jason Furman and Peter Orszag. A firm-level perspective on the role of rents in the rise in inequality. *Working Paper*, 2015.

Andrew Glover and Jacob Short. The age-distribution of earnings and the decline in labor's share. *Working Paper*, 2016.

Douglas Gollin. Getting income shares right. *Journal of Political Economy*, 110(2):458–474, April 2002.

Gustavo Grullon, Yelena Larkin, and Roni Michaely. Are US industries becoming more concentrated? *Working Paper*, 2016.

Barney Hartman-Glaser, Hanno N. Lustig, and Mindy X. Zhang. Capital share dynamics when firms insure managers. *Working Paper*, 2017.

Cosmin Ilut, Matthias Kehrig, and Martin Schneider. Slow to hire, quick to fire: Employment dynamics with asymmetric responses to news. *Journal of Political Economy*, forthcoming.

Nir Jaimovich and Henry Siu. The trend is the cycle: Job polarization and jobless recoveries. *Working Paper*, 2015.

Nicholas Kaldor. Capital accumulation and economic growth. In F. A. Lutz and D. C. Hague, editors, *Theory of Capital*, pages 177–222. St. Martin's Press, New York, 1961.

Loukas Karabarbounis and Brent Neiman. The global decline of the labor share. *Quarterly Journal of Economics*, 129(1):61–103, February 2014a.

Loukas Karabarbounis and Brent Neiman. Capital depreciation and labor shares around the world: Measurement and implications. *Working Paper*, 2014b.

Matthias Kehrig. The cyclicality of productivity dispersion. *US Census Bureau Center for Economic Studies Paper No. CES-WP-11-15*, May 2011.

Dongya Koh, Raül Santaeulàlia-Llopis, and Yu Zheng. Labor share decline and intellectual property products capital. *Working Paper*, 2016.

Robert Z. Lawrence. Recent declines in labor's share in US income: A preliminary neoclassical account. *NBER Working Paper No. 21296*, 2015.

Marc Melitz and Saso Polanec. Dynamic Olley-Pakes productivity decomposition with entry and exit. *RAND Journal of Economics*, 46(2):362–375, Summer 2015.

Marc J. Melitz and Gianmarco I. P. Ottaviano. Market size, trade, and productivity. *Review of Economic Studies*, 75(1):295–316, January 2008.

Ezra Oberfield and Devesh Raval. Micro data and macro technology. *NBER Working Paper No. 20452*, 2014.

G. Steven Olley and Ariel Pakes. The dynamics of productivity in the telecommunications equipment industry. *Econometrica*, 64(6):1263–1297, November 1996.

José-Víctor Ríos-Rull and Raül Santaeulàlia-Llopis. Redistributive shocks and productivity shocks. *Journal of Monetary Economics*, 57(8):931–948, November 2010.

Matthew Rognlie. Deciphering the fall and rise in the net capital share: Accumulation or scarcity? *Brookings Papers on Economic Activity*, Spring 2015.