THE MICRO-LEVEL ANATOMY OF THE LABOR SHARE DECLINE

Matthias Kehrig
Nicolas Vincent

The Micro-Level Anatomy of the Labor Share Decline
Matthias Kehrig and Nicolas Vincent
NBER Working Paper No. 25275
November 2018
JEL No. E2,L1,L20,L6,O4

ABSTRACT

The aggregate labor share in U.S. manufacturing declined from 62 percentage points (ppt) in 1967 to 41 ppt in 2012. The labor share of the typical U.S. manufacturing establishment, in contrast, rose by over 3 ppt during the same period. Using micro-level data, we document a number of striking facts:(1) there has been a dramatic reallocation of value added to "hyper-productive" (HP) low-labor share establishments, with much more limited reallocation of inputs; (2) HP establishments have only a temporarily lower labor share that rebounds after five to eight years to the level of their peers; (3) selection into HP status has become increasingly correlated with past size; (4) labor share dynamics are driven by revenue total factor labor productivity, not wages or capital intensity; (5) employment has become less responsive to positive technology shocks over time; and (6) HP establishments enjoy a product price premium relative to their peers that causes their high (revenue) productivity. Counterfactual exercises indicate that selection along size rather than shocks or responsiveness to them is the primary driver of the labor share decline.

Matthias Kehrig
Department of Economics
Duke University
237 Social Sciences, Box 90097
Durham, NC 27708
and NBER
matthias.kehrig@duke.edu

Nicolas Vincent
HEC Montreal
Institute of Applied Economics
3000, chemin de la Côte-Sainte-Catherine
Montréal (Québec), 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 important for a number of reasons. For one, 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 suggests 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.

The literature has proposed numerous explanations for that aggregate decline, most of which are rooted in firm-level behavior. But little is known about the empirical dynamics at the micro level. Our paper fills this gap. We use confidential data from the U.S. Census of Manufactures to study the establishment- and firm-level anatomy of the labor share in manufacturing, a sector which is overwhelmingly responsible for the decline of the labor share in the entire private economy.

We document a number of striking facts, some of them difficult to reconcile with the multiple channels that have been suggested by various authors. First, we confirm that the labor share in the manufacturing sector declines by almost 5 percentage points (ppt) per decade between 1967 and 2012. This, however, hides contrasting dynamics at the micro level: Alongside the aggregate decline, the median establishment saw an *increase* in its labor share, by about 0.7 ppt per decade. In fact, this upward trend is present for the vast majority of manufacturing establishments. We find that the decline of the aggregate labor share is entirely driven by a strong reallocation of value added to establishments with low labor shares (see Figure 1). In contrast, reallocation of labor, materials and capital was much less pronounced over the same period.

Second, we find that labor shares at the establishment level are strikingly transient. We define "hyper-productive" (HP) establishments as those in the lowest quintile of the labor share distribution for a given year and sector. The probability that a typical HP establishment today loses that status five years later is about 60%, while that number would be close to 0% if HP establishments had permanently low labor shares. Even more strikingly, we document that the labor share dynamics of HP establishments follow a V-shape pattern: the drop in labor share they experience in the five years preceding HP status is almost equal to the rebound in the following five years. In addition, the depth of the V-shape is found to be increasing over time meaning that HP and Non-HP establishments look increasingly different.

Third, we document the presence of a significant size selection effect in the data: the larger is an establishment today, the more likely it will turn out as a HP (low labor share) establishment in the next Census year. This phenomenon has become much more pronounced over time to the point that, by 2007, an establishment in the largest size quintile was three to four times more likely to become a HP establishment five years later than an establishment in the bottom quintile.

Fourth, we decompose the labor share into its components to understand the main drivers behind the micro-level dynamics. We find that the V-shaped pattern of the labor share of HP establishments is due almost entirely to fluctuations in revenue total factor productivity (TFPR),

Figure 1: The changing distributions of labor shares and value added
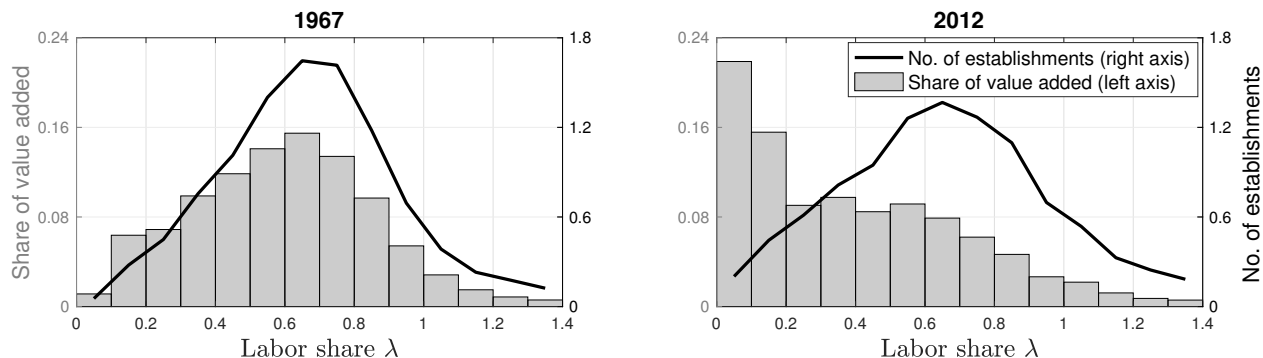
Note: The solid black line (right axis) reflects the raw cross-establishment distribution of labor shares. These pure numbers of establishments show no significant locational shift of establishment-level labor shares from 1967 to 2012; the fattening of tails indicates a polarization of labor shares that does not affect the aggregate labor share by itself. The distribution of economic activity (value added shares in grey bars, left axis), in contrast, dramatically shifts towards low-labor share establishments. This reallocation of value added is principally responsible for the aggregate labor share decline.

To account for industry-specific differences in the raw and value added-weighted labor share distributions, they are first calculated within each 3-digit NAICS industry. Then these distributions are averaged across these 21 manufacturing industries using value added weights in a given year to obtain an estimate of the typical within-industry distribution of raw and value added labor shares in that year. Table 4 confirms that between-industry reallocation plays only a minor role for the aggregate labor share decline.

with very little accounted for by wage or capital intensity variations.

Fifth, we show that there has been a rising disconnect between employment growth and TFPR shocks in our sample over time. In the 1970s, when the aggregate labor share was stable, employment used to respond symmetrically to negative and positive productivity shocks. By the 2000s, at a time when the labor share was declining strongly, those establishments experiencing positive TFPR shocks, such as HP plants, showed instead no more inclination to hire than their peers.

Sixth, we use a subsample of the Census database which provides information about the value of sales and quantity for individual products. This allows us to derive the contribution of the "product price premium" (an establishment's deviation from the average price of its competitors) to differences in sales per worker across establishments and over time. We find that low-labor-share establishments tend to have on average significantly higher prices than their peers, and that the dynamics of the price premium is first-order in understanding the dynamics of sales per worker of HP establishments.

Next, we construct counterfactual scenarios in order to determine what are the micro-level factors that are most likely behind the decline in the aggregate labor share. To do so, we focus on three candidate explanations related to time-series facts we documented earlier: the increasing depth of the V-shaped labor share pattern of HP establishments, the declining responsiveness of employment to positive value-added shocks and the rising correlation between past size and HP status. We find that the latter plays an important role: keeping the selection effect constant to what it was at the beginning of the sample eliminates between 65% and 85% of the observed aggregate

labor share decline.

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) and Acemoglu and Restrepo (2018) refine this theory by focusing on information and communication technology capital or robots, respectively. Koh, Santaeulàlia-Llopis, and Zheng (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. 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. Some empirical work by Lawrence (2015) and Oberfield and Raval (2014) casts doubt on that even at high levels of aggregation. But even if capital and labor are complements, Grossman, Helpman, Oberfield, and Sampson (2017) show that slowing growth in labor- or capital-augmenting technological change can lead to a labor share decline. Alvarez-Cuadrado, Long, and Poschke (2015) show that industry-level specificities in technological change and the elasticity of substitution between capital and labor matter for the dynamics of industry-level factor shares.

Alternatively, Elsby, Hobijn, and Şahin (2013) advocate the role of offshoring as an important driver of the labor share decline. In related work, Boehm, Flaaen, and Pandalai-Nayar (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 (2018) find the age composition of the workforce has shifted towards workers that are less capable of extracting their marginal product of labor as a wage. Kaymak and Schott (2018) document a relationship between cuts in corporate tax rates and labor share declines in OECD countries.

Furman and Orszag (2015) noted that the distribution of capital returns – inversely related to the labor share – had shifted up and became more skewed towards high-return firms. Hartman-Glaser, Lustig, and Zhang (forthcoming) 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's owner can ask for a larger insurance premium, raising in turn the capital share. This is consistent with the finding in Kehrig (2011) that the productivity dispersion across establishments has increased significantly. 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, Dorn, Katz, Patterson, and Reenen (2017a,b), for example, present some industry- and establishment-level evidence on firm concentration shares which is consistent with our finding that a small fraction of "hyper-productive establishments" are mainly responsible for the aggregate labor share decline. Grullon, Larkin, and Michaely (2016) use firm-level data

4

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) and Eeckhout and De Loecker (2017) show that markups have grown over time, lowering both the labor and capital shares. Edmond, Midrigan, and Xu (2018) finally study the welfare implications of high markups and high markup dispersion. They find that reducing markups by taxing large high-markup firms may reduce concentration but also welfare. Like us, they carefully examine the demand side sources of profitability and labor share dynamics. Baqaee and Farhi (2017) study misallocation in networks and find that high-markup firms have gotten larger over time which is consistent with our finding that few but large low-labor share establishments generate very high revenue labor productivity. This is also corroborated by findings in Neiman and Vavra (2018) who use household scanner data to show that firms are increasingly able to introduce customized products that make up a large share of individual consumer spending.

Our finding of lots of turnover among highly productive low labor share units is consistent with the findings in Brynjolfsson, McAfee, Sorell, and Zhu (2008). They establish that IT investment enables better scalability thus making it possible for individual firms to quickly generate large sales that we observe in the data. They also find that those IT intensive industries are typically more concentrated and exhibit higher turnover.

Issues related to the measurement of the labor share abound: Elsby, Hobijn, and Şahin (2013) refine the imputation of the labor portion of noncorporate income, an adjustment that only moderately mitigates the labor share decline. 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. manufacturing labor share

Our focus is on the labor share dynamics in the U.S. manufacturing sector. Despite its declining economic weight, we choose to concentrate our attention on this sector for a number of reasons. First, as highlighted by Elsby, Hobijn, and Şahin (2013), manufacturing is one of the sectors where the labor share decline has been most pronounced. In fact, it accounts for almost all the fall in the labor share for the private non-farm economy. This makes it a natural starting point to study the macro and micro dynamics of the labor share decline. Second, many of the explanations commonly put forward to explain the fall in the labor share, such as automation, competitive pressures by globalization, offshoring, the fall in the power of labor unions, etc. are particularly relevant in the context of goods-producing activities.

Third, data at the level of individual manufacturing establishments from the U.S. Census Bureau have been heavily studied and are considered to be of higher quality than for other sectors.

For example, the information on intermediate inputs and energy use contained in the Census of Manufactures database allows us to construct reliable measures of value added, instead of having to rely on alternative variables such as sales or revenue to generate establishment-level labor shares. Fourth, the longer time coverage for the manufacturing sector makes it possible to contrast the dynamics of the labor share both before and after the start of its secular decline, around the early 1980s. While our analysis starts in 1967, the U.S. Census Bureau only began to sample establishments in other sectors in the 1980s or 1990s. Fifth, the higher degree of homogeneity for some manufacturing goods will allow us to disentangle the respective roles of prices and quantities in driving the phenomena we document in the following sections. Finally, 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, Hobijn, and Şahin (2013) argue that self-employment income matters significantly for these trends. In addition, our results are unlikely to be biased by the evolution of housing prices that impact household-level surveys: 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. Finally, computations by Koh, Santaeulàlia-Llopis, and Zheng (2016) show that manufacturing is one of the few sectors, in which the labor share decline is not overturned by the rise in intellectual property products.

2.1 Data sources and measurement

Most of the results derived throughout the paper come from the establishment-level Census of Manufactures database. The U.S. Census Bureau collects data on all manufacturing establishments within the Economic Census, which is taken every five years from 1967[1] until 2012. We drop all observations that are administrative records or are not part of the "tabbed sample" which makes up the official tabulations published by Census.

For comparison purposes, we also rely on industry-level data from the Bureau of Labor Statistics (BLS). This dataset comes from the annual "KLEMS Multifactor Productivity Tables by Industry" for both Manufacturing and Nonmanufacturing Industries and spans 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.

In either dataset, the labor share λ_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. Focusing on the raw nominal data has the advantage of avoiding measurement issues related to inflation.

---

[1]The 1963 Census lacks a substantial portion of labor compensation, so we ignore it in this paper.

In the BLS data, labor costs comprise employee compensation (wages, salaries and supplements) as well as a portion of non-corporate income.[2] 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 `ILC`) such as unemployment insurance or social security contributions netted out from wages and voluntary labor costs (item `VLC`) such as health, retirement and other benefits paid to employees.[3] Value added in the Census data is measured as sales less inventory investment for final and work-in-progress goods, resales, material inputs and energy expenditures.[4] In addition, we drop all observations in the bottom and top percentiles to avoid that outliers drive our results. This implies that we discard observations with a negative value added (and thus labor share).[5]

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 larger. Third, the approach we use to adjust for 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 that computed from BLS data.

2.2 The labor share decline in U.S. manufacturing

Before going any further, 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 establishments 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. We also include a series for

---

[2]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).

[3]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.

[4]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.

[5]For more details on the construction of the sample and the variables of interest, see Appendix A.

the non-farm private sector excluding manufacturing.

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

Note: The solid black line (left scale) represents the aggregate labor share λ_t in the Census of Manufactures panel as calculated in Equation (2); the thin grey line with balls represents the aggregate labor share in the manufacturing sector as calculated from BLS data. The aggregate labor share in the non-manufacturing sectors is displayed as the solid grey line and remains largely constant.

Figure 2 confirms that the BLS labor share is about 8 ppt higher on average, for the reasons stated above. Yet, the two labor series exhibit very similar 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 4.5 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 6.2 ppt per decade since the 1982 Census.

Lastly, while the aggregate labor share declines strongly in manufacturing, that downward trend is essentially absent for the non-manufacturing sectors of the private economy. Their labor share fluctuates around 67 ppt until the late 1990s, increases somewhat before declining to settle just above 65 ppt. The implication is that the overall labor share decline in the private non-farm economy appears to be largely driven by manufacturing, which rationalizes our focus on this sector.

3   Six findings about the labor share

In this section, we study the anatomy of the decline in the aggregate labor share by exploiting establishment-level data. We present and analyze six main findings on the micro-level dynamics of the labor share.

3.1 Finding 1: The labor share – aggregate decline, micro-level increase

As a starting point, we rely on the folllowing decomposition of the aggregate labor share λ_t:

$$\lambda_t = \frac{\sum_i W_{it} L_{it}}{\sum_i Y_{it}} = \sum_i \lambda_{it} \omega_{it} \tag{2}$$

where λ_{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. Our aim at this point is to study the respective roles of the distribution of individual labor shares and reallocation across establishments in driving down the aggregate labor share.

3.1.1 The labor share of the median establishment *increases*

As a first step, Figure 3 plots several quantiles of the raw distribution of establishment-level labor shares λ_{it} in each Census year since 1967, alongside the aggregate labor share.

Figure 3: Aggregate and establishment-level labor shares

Note: The figure plots the aggregate labor share (black line, left axis) against the year-by-year quantiles of the cross-establishment labor share distribution (grey lines, right axis): the solid grey line with balls reflects the median, the dashed grey lines reflect the first and third quartile. While the aggregate labor share declines strongly, the median and top quartile labor share increase over time.

Figure 3 highlights diverging trends in the labor shares at the aggregate and establishment level, particularly since the mid 1980s: while the aggregate labor share declines by 4.5 ppt per decade on average, the median labor share *increases* by 0.7 ppt per decade. The top and bottom quartiles strongly co-move with the median and increase as well. An implication of this finding is that the aggregate labor share decline is not the result of a shift of the distribution of labor shares in individual establishments (corresponding to the λ_{it} terms in Equation (2)). Instead, our evidence points to the importance of reallocation (corresponding to the ω_{it} terms in Equation (2)) as the main driver of the aggregate labor share dynamics. This is what we turn our attention to next.

9

3.1.2 The dramatic reallocation of value added

So far, we have considered changes in the distribution of establishment-level labor shares, corresponding to the λ_{it} terms in the decomposition of the aggregate labor share of Equation (2). In the previous section we found that the upward trend of the median establishment-level labor share stood in stark contrast to the decline of the aggregate labor share over the 45 years covered by our sample. These diverging trends imply that the ω_{it} terms in Equation (2) 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. To quantify this reallocation margin, we divide the distribution of labor shares λ into 10 ppt-wide bins from 0 to 140 ppt in each year. We then compute the share of aggregate value added as well as the share of establishments accounted for by each labor-share bin. To control for industry-specific differences, we compute these shares for each 3-digit NAICS industries and then aggregate them up in each bin using the industry's value-added weight in the given year. The subsequent analysis of reallocation of value added therefore refers to reallocation *within* a typical industry.

Figure 4: Value-added weights and labor share distribution
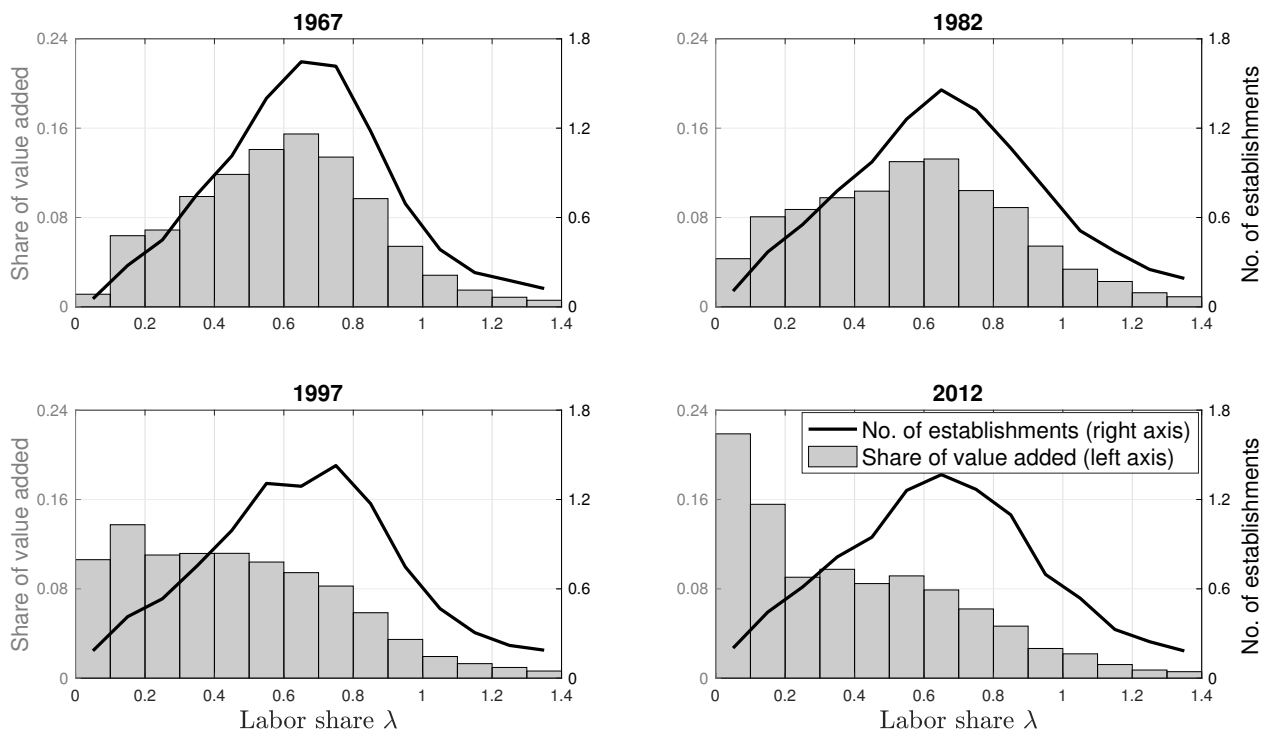
Note: See notes to Figure 1.

   Figure 4 jointly displays the distributions of labor shares λ_{it} and value-added weights ω_{it} every 15 years, from 1967 to 2012. The panels paint a striking picture: most of value added in 1967 is produced by establishments with a middle-of-the-road labor share (between 50 and 80 ppt). The value-added weighted median labor share is 62 ppt. Over the following decades, however, the

economic activity shifts gradually and persistently to the low-labor share spectrum. By 2012, half of aggregate value added if accounted for by establishments with a labor share less than 32 ppt. This evidence has been confirmed for other sectors in the U.S. by Autor, Dorn, Katz, Patterson, and Reenen (2017b), Gouin-Bonenfant (2018) for Canada and Berkowitz, May, and Nishioka (2017) for China, who also document a massive reallocation of output over time. Next, we investigate whether this phenomenon was driven by entry and/or exit, and whether it is matched by a similar reallocation of inputs.

3.1.3   The limited role of entry and exit

The dramatic reallocation of value added documented in Figure 4 could be happening through entry and exit (extensive margin) and/or differential value added growth rates of incumbent establishments (intensive margin). The extensive margin is quantatively relevant if, in a given year, exiting establishments tend on average to have labor shares that are significantly above those of incumbents and/or entrants have relatively low labor shares. To assess the importance of the extensive margin, we compute the aggregate labor share for a strongly balanced sample of establishments that are permanently active from 1967 to 2012.

Figure 5: Aggregate labor share in full and strongly balanced sample

Note: The figure plots the aggregate labor shares computed on the full panel (solid black line) against that computed on a strongly balanced panel (solid grey line). It shows that entry and exit matter for the *level*, but not for the *decline* of the labor share.

Figure 5 reveals that entry and exit do have an impact on the level, but not the downward trend of the aggregate labor share. The aggregate labor share in the strongly balanced sample is about 2 ppt lower than that of the full sample, suggesting that entry and exit indeed depresses the labor share. Yet, except for the last year in the sample, the labor share declines in the full and balanced samples are virtually indistinguishable: the aggregate labor share in both samples stagnates until 1982; from 1982 to 2007 it falls by 7.3 ppt per decade in the full sample and by 7.4 ppt in the balanced panel. Our conclusion is that the contribution of the extensive margin to the

decline of the aggregate labor share is quantitatively negligible. Instead, most reallocation leading to the secular decline takes place among incumbent establishments.

3.1.4   The essential role of the bottom tail

In subsequent sections, we will often focus on the characteristics of establishments with the lowest labor shares. This focus is particularly relevant if most of the dramatic reallocation documented above takes place at the bottom of the labor share distribution, instead of throughout. This is what we turn our attention to next.

Defining "hyper-productive" (HP) establishments   In order to contrast the dynamics of low-labor share establishments relative to those of their peers, we define establishments in the lowest quintile of the labor share distribution in a given year and 3-digit NAICS industry as "hyper-productive," abbreviated as "HP establishments." The quintiles are industry-specific due to the wide range of average labor shares across industries. As we have seen before, the value-added weight of these HP establishments has risen dramatically over time relative to their Non-HP peers.

HP establishments are the key driver of the aggregate labor share   To highlight the role of HP establishments in shaping the dynamics of the manufacturing sector, we re-compute the aggregate labor share by simply dropping the HP establishments. If reallocation was pervasive throughout the distribution, we would expect to also observe a labor share decline in this subsample, albeit from a higher starting point.

Figure 6: The importance of HP establishments

Note: The figure plots the aggregate labor shares computed on the full panel (solid black line) against that computed for the panel after dropping the set of HP establishments (solid grey line). It shows that Non-HP establishments do not contribute to the decline of the aggregate labor share.

The labor shares including and excluding HP establishments are shown in Figure 6. Two aspects stand out: First, not surprisingly, the level of the aggregate labor share is much higher at about 75 ppt. Second, and more importantly, it does not exhibit any significant decline: while the

actual aggregate labor share starts to fall in the 1980s, the counterfactual aggregate labor share without HP establishments continues to fluctuate around 75 ppt. In other words, while reallocation among Non-HP establishments may be taking place, it does not contribute meaningfully to the empirically observed aggregate labor share decline.

The takeaway from this exercise is clear: in order to understand the aggregate labor share decline, it is essential to study the characteristics of HP establishments, those at the bottom of the labor share distribution.

3.1.5 Weak reallocation of inputs

Is this massive reallocation of value added mirrored by a reallocation of the inputs used by establishments? Standard firm dynamics models such as Jovanovic (1982) and Hopenhayn (1992) posit that large, highly productive firms should also account for a large portion of input use. Knowing whether or not both output and inputs get reallocated towards low-labor share establishments may shed light on the drivers of the aggregate labor share decline. For example, one possibility is that low-labor share establishments become increasingly capital intensive, another would be that they do not accumulate significantly more capital, but enjoy large gains in total factor productivity without hiring more workers and paying higher wages.

Figure 7 displays the reallocation of economic activity from 1967 to 2012. The top-left panel repeats the evidence for value added from Figure 4, while the other panels present the analogue for the main inputs: capital goods (top right panel), hours worked (bottom left panel) and materials used (bottom right panel). Overall, the figure shows that the reallocation of inputs is much more limited when compared to the strong reallocation of value added towards the low end of the labor share distribution. The contrast is remarkable: In 1967, establishments with a labor share of 20 ppt or less produce only 7.5% of aggregate value added with about 3% of labor input and 7% of the capital stock. By 2012, establishments in that portion of the labor share distribution now see their combined value added make up 37.5% of the aggregate, yet account for a mere 10.6% of labor input and operate 22% of capital.

During the same period, the value-added-weighted median drops from 61.2% to 34.5%. This means that half of manufacturing value added in 2012 is produced by establishments with a labor share of 61.2% or less (34.5% in 1967). The analogous numbers for capital and labor inputs are much less pronounced: While half of the capital stock in 1967 is employed by establishments with a labor share of 60.7% or less, this number merely drops to 47.4% in 2012. For labor input, the reallocation is even less pronounced: the median labor share weighted by hours worked declines from 71.5% to 66.9%.[6] To summarize: the value added-weighted median labor share drops more than twice as much as the capital-weighted median labor share while the labor-weighted median labor share essentially stagnates. In other words, both nominal labor and capital productivity for these establishments greatly increased over time.

---

[6]Repeating this exercise with other measures (total employment, production workers and production hours) yields similar results.

Figure 7: Dramatic reallocation of output, limited reallocation of inputs
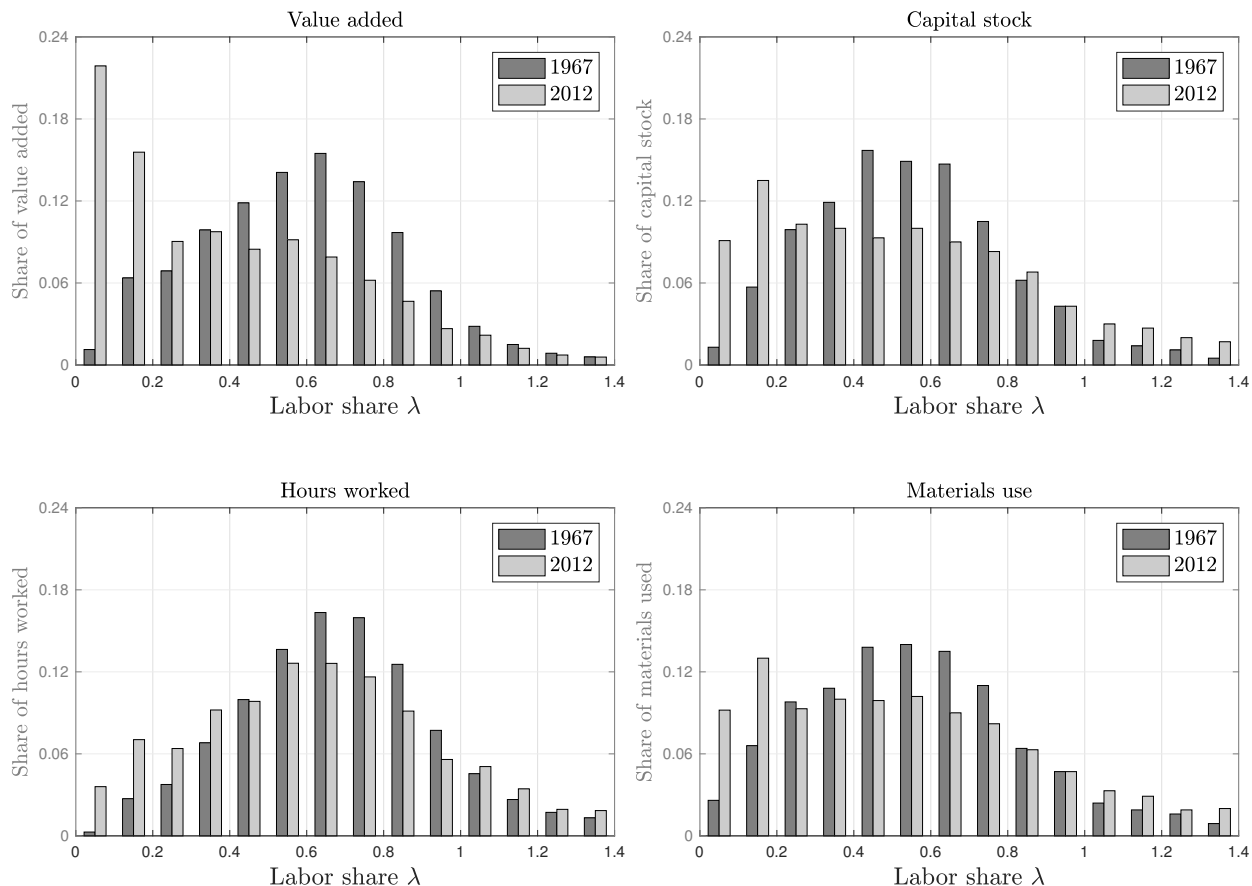
Note: Top left panel depicts the distribution of value added across the labor share spectrum (the same as the top left and bottom right panels of Figure 4) in 1967 versus 2012. The remaining panels depict the analogue of the 1967-vs-2012 allocations of capital (top right), hours worked (bottom left) and materials use (bottom right).

3.2 Finding 2: Low labor shares are transient

So far, our analysis of the labor share decline was static in nature. We focused on cross-sectional snapshots of the data, analyzing the distributions of labor shares, value added and inputs across establishments in each Census year. While this approach is common in the literature, it turns a blind eye to the transitional dynamics happening *within* the distribution. These dynamics, in turn, can teach us about the forces and factors that lie behind the decline in the labor share: the nature of shocks hitting individual establishments, the technologies they operate and the frictions they face.

3.2.1 Markov transitional dynamics

We start by documenting the transition dynamics of HP and Non-HP establishments with the help of a simple Markov transition matrix, displayed in Table 1. That is, we ask a simple question: conditional on an establishment's labor share at time t, what is the probability that it has HP status

at time $t + 5$? If establishment-level labor shares were highly persistent and HP establishments were to retain this status, this probability should be equal to 100%. On the polar opposite, in an economy where an establishment's labor share was drawn randomly every year, this number would be 20% (one fifth as we define HP in terms of the lowest quintile).

Table 1: Transition probabilities of HP status

Panel A. Unweighted transitional dynamics

	Non-HP_{t+5}	HP_{t+5}
Non-HP_t	0.854	0.146
HP_t	0.583	0.417

Panel B. Weighted transitional dynamics

	Non-HP_{t+5}	HP_{t+5}
Non-HP_t	0.922	0.078
HP_t	0.536	0.464

Note: Markov matrix of labor shares from Census to Census. Panel A. considers the share of establishments that remain/leave/enter HP status when quintiles are unweighted, Panel B. displays the share of aggregate value added accounted for by the HP establishments when defined by VA-weighted quintiles.

Table 1 shows that over our sample period, the probability that an establishment retains HP status from Census year to Census year (a 5-year window) is only 41.7%. While this is higher than if HP status were perfectly random (20%), the transition probability indicates that labor share at the establishment level is surprisingly transient, even for the most productive establishments.

The evidence in Table 1 was obtained by counting each establishment as one unit, irrespective of its size. One may be concerned that these results are mostly driven by small, economically insignificant establishments. To account for dynamics of economic activity, we also consider Markov transition matrices of quintiles weighted by economic activity and confirm the transient dynamics of HP establishments. These results are displayed in Panel B.; they indicate that there is slightly more persistence when considering transitional dynamics weighted by value added, but the overall impression of little persistence remains.

3.2.2 Labor shares of HP establishments follow a V-shaped pattern over time

We now investigate more precisely the relative labor share dynamics of HP and Non-HP establishments. First, we construct backward-looking (from years $t - 5$ to t) and forward-looking (from t to $t + 5$) percentage-point changes in establishment labor shares from the Census panel. We then regress these constructs on a dummy variable that equals one if an establishment is among the HP establishments in the current Census year:

$$d\lambda_{it} \equiv \lambda_{it} - \lambda_{it-5} = c_1 + \beta_{-5}\mathbb{I}\{HP_{it}\} + \gamma_1 X_{it} + \varepsilon_{1it} \tag{3}$$

$$d\lambda_{it+5} \equiv \lambda_{it+5} - \lambda_{it} = c_2 + \beta_{+5}\mathbb{I}\{HP_{it}\} + \gamma_2 X_{it} + \varepsilon_{2it} \tag{4}$$

15

While the *level* of the labor share of HP establishments is below that of their peers by definition – they consist of all establishments in the lowest quintile in a given year and industry –, our aim here is to analyze their *dynamics* from the estimates of the coefficients β_{-5} and β_{+5} in Equations (3) and (4). That is, we study how the labor share dynamics of HP establishments differ from those of Non-HP establishments over a ten-year window around the reference period. Note that we do not require that HP establishments in period t were also HP establishments in $t-5$ and will be in $t+5$; an establishment could well have HP status for a single year. The vector X_{it} contains industry, region and year dummies as controls. We estimate Equations (3) and (4) once using the unweighted Census panel and again using value-added weights to account for the fact that larger establishments likely have less volatile labor shares. Results are displayed in Table 2.

Table 2: The dynamics of HP establishments

Variable	(I)	(II)	(III)	(IV)
β_{-5}	−0.2883***		−0.1755***	
	(0.0068)		(0.0083)	
β_{+5}		+0.2717***		+0.1491***
		(0.0069)		(0.0084)
R^2	0.111	0.096	0.102	0.070
Weights	none	none	VA weights	VA weights

Note: Pooled OLS regression of Equations (3) and (4) on the full Census panel. "VA weights" correspond to the share of establishment i's value added in aggregate value added. Standard errors are clustered at the 4-digit NAICS industry level. Significance levels are denoted by * (10% level), ** (5% level), and *** (1% level).

If the labor share of HP establishments were permanently low, the coefficients β_{-5} and β_{+5} would be small (in absolute value), indicating little difference in labor share dynamics between HP and Non-HP establishments. The estimation shows that this is clearly not the case: β_{-5} is estimated to be strongly negative while β_{+5} is positive. Both coefficients are statistically significant at the 1% level.

The coefficient estimate from the unweighted regression in column (I) implies that relative to the previous Census year, an establishment that has HP status at time t saw a change in its labor share over that 5-year span that is 28.8 ppt lower than that of its Non-HP peers. In the five-year period thereafter, the coefficient estimate of β_{+5} in column (II) indicates that the change in the labor share of establishments that are HP in year t will be 27.2 ppt *higher* than that of Non-HP establishments. Columns (III) and (IV) in Table 2 repeat the same regression while weighing establishments by their value-added share; the broad findings are similar, even though the magnitudes are smaller (−17.6 ppt and +14.9 ppt, respectively). To account for economic activity, all subsequent dynamic analyses will be weighted by value added unless otherwise noted.

To ease the interpretation, we also report the results for β_{-5} and β_{+5} as cumulative growth rates in the left panel of Figure 8. It is striking to see that while the relative labor share change of time-t HP establishments is very negative between $t-5$ and t, by $t+5$ these HP establishments appear

to be no more different than their Non-HP peers than they were 10 years earlier. All in all, our analysis appears to show that the average HP establishment experiences a rather temporary drop and rebound in its labor share, rather than remaining at a permanently or even highly persistent lower level. This is in line with the earlier evidence from the Markov transition matrices.

Figure 8: The temporary fall and rise of labor shares of HP establishments

Note: Left panel: Cumulative evolution of the labor share of the average HP establishment relative to their peers in the Census panel before ($t-5$ to t) and after (t to $t+5$) the year it is in HP status. Unweighted dynamics in dark grey, value added-weighted dynamics in light grey; whiskers denotes 95% error bands.
Right panel: Analogous labor share dynamics of HP establishments in the ASM data.

To ease the interpretation, in the left panel Figure 8 we report the results for β_{-5} and β_{+5} as cumulative growth rates.

Before moving on, it should be noted that while Figure 8 makes clear that both unweighted and value-added-weighted estimates of HP labor share dynamics exhibit a V-shaped pattern, there are some HP establishments that do display permanently low labor shares. Although these establishments tend to be larger than the average, an exercise similar to the one in Section 3.1.4 shows that they are responsible for only a fraction of the aggregate labor share decline (see Appendix B.1).

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 establishments would simply be establishments that experienced huge mismeasurement at time t, yet whose fundamentals were not any different than the typical establishment in the population. This would mechanically give rise to the temporary change shown in the left panel of Figure 8. Using only data every five years would make our analysis vulnerable to measurement errors in just that single year.

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 establishments, 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 establishments that are HP establishments

17

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).

For this robustness check, we adapt the estimation in Equations (3) and (4) to an annual frequency and run ten regressions, one for each of the preceding five and subsequent five years. The results are reported in the right panel of Figure 8: They confirm the transient nature of the labor share that we found using the Census years. However, while the trough at t is unmistakable, notice that the relative change in the labor share is not taking place entirely between $t-1$ and t, but instead regularly over the preceding years. Also, notice that it does not recover fully even after five years, when the labor share is estimated to still be 5 to 8 ppt below the level of Non-HP establishments. All in all, our evidence appears to indicate that the transient nature of HP status is not merely an artifact of measurement error.[7]

We also consider the possibility that the transient pattern of establishment-level labor shares merely reflects reallocation dynamics within firms. To that end, we aggregate establishments to the level of the firm (a firm in the Census is defined as the organization that has organizational control over establishments rather than the firm as the employer identification number, EIN) and define HP firms analogously to our establishment-level definition. The results are displayed in Appendix B.2 and are largely similar to the ones we discussed in this section.

3.2.3 The evolution of the V-shaped pattern over time

The results so far were obtained by pooling all years together. Figure 9 illustrates the evolution over time of the V-shaped pattern that we documented earlier. More specifically, for each Census year t we compute the change in the labor share between $t-5$ and t for time-t HP $(d\lambda_t^{HP})$ and Non-HP $(d\lambda_t^{\text{Non-}HP})$ establishments, and plot the difference. Both unweighted and value-added-weighted estimates are shown. In both cases, the evidence is striking: starting from the mid-1980s, the labor share dynamics of HP and Non-HP get increasingly different, with an unweighted differential of 21 ppt in 1972 that increases to 38 ppt by 2012. This increase in the depth of the V-shape pattern is even more pronounced in relative terms when observations are value-added-weighted (from 12 ppt to 28 ppt).

3.3 Finding 3: Low labor shares are increasingly correlated with past size

Since the aggregate labor share depends not only on the establishment-level labor share distribution but also the joint labor-share/size distribution, the *types* of establishments that are selected to become HP is likely to matter. In this section, we show that selection into HP status is a positive

---

[7]As an additional way to support our finding of the transient nature of low labor shares, we turn to the "product trailer" of the CMF. This portion of the Census records sales for individual products of an establishment. We use these product sales numbers, sum them up to the establishment-year level and thus obtain an alternative sales/labor share measure. The labor share dynamics of HP establishments using this alternative sales measure are very similar to our benchmark presented in Figure 8. In that exercise, we drop sales recorded under product balancing codes and omit imputed product values to guard against the problems associated with imputation highlighted in White, Reiter, and Petrin (2018), so the sum of product sales and total value of shipments do not necessarily coincide.

Figure 9: Labor share change of HP versus Non-HP establishments over time

Note: This figure displays the difference in labor share dynamics between HP and Non-HP establishments (corresponding to the $t-5$ to t bars in the left panel of Fig. 8) by year. It shows that HP establishments look increasingly different from their peers starting in 1987.

function of establishment size, and that this relationship has stenghtened over time.

We run the following specification 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, while X_{it} is a vector of industry and regional dummies. Each β coefficient captures the probability that a time-$t+5$ HP establishment came from a specific size quintile at time t. If the size of an establishment 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.

The two lines in Figure (10) report the ratio β_5/β_1 from Regression (5), both weighted and unweighted. The fact that this ratio is always above 1 indicates that on average, larger establishments (those in the fifth quintile) are more likely to have HP status five years down the road.

More strikingly, this relationship has greatly strengthened over the years: 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, the relative probability in the unweighted regression was about 1.3 in 1967, but this relative probability more than doubled by 2007. As can be seen from Figure 10, both the level and change in the relative probability are even more pronounced in a value added-weighted regression.

Figure 10: Relative probability of becoming HP: largest vs. smallest size quintile

Note: Figure reports ratio of β_5/β_1 from running regression in Equation (5), run separately by year=1967, ..., 2007. Size quintiles are defined along establishment 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).

3.4 Finding 4: Low labor shares are driven by high TFPR, not high capital intensity or low wages

In the previous sections, we documented that low-labor share establishments have accounted for an increasingly large portion of aggregate value added; that HP status is remarkably transient; and that size has become a better predictor of future HP status with time. We now turn our attention to the components of the labor share.

The log of the labor share of establishment i at time t can simply be written as

$$\log \lambda_{it} = \log W_{it} - \log(Y_{it}/L_{it}) \tag{6}$$

$$= \log W_{it} - \alpha \log(K_{it}/L_{it}) - \log TFPR_{it} \tag{7}$$

where on the first line, W_{it} is the wage of the average employee and Y_{it}/L_{it} is revenue labor productivity.[8] Assuming that firms operate a constant-returns-to-scale Cobb-Douglas production function, we can split labor productivity in the second line into two subcomponents: (scaled) capital intensity, $\alpha \log(K_{it}/L_{it})$ with α being the production elasticity of capital, and nominal productivity, $\log TFPR_{it}$.

To ensure that our results are not driven by systematic cross-industry differences as well as to make them more readily interpretable – wages, value added per worker and capital intensity are nominal variables –, we study an establishment's wage, capital intensity and TFPR relative to that

---

[8]It is important to notice that wages and labor productivity are both nominal variables. In addition, $\log(Y_{it}/L_{it})$ is nominal productivity, that is, it compounds both physical labor productivity and prices. In the language of the recent productivity literature, we study *revenue* labor productivity. In the next section, we will differentiate between revenue labor productivity and physical labor productivity, the analogue of TFPQ in Foster, Haltiwanger, and Syverson (2008); Hsieh and Klenow (2009).

of its peer group. We define peers to be establishments that are active in the same product and labor markets: for a given year, a peer group corresponds to all establishments in the same state and 3-digit NAICS industry.[9] The relative wage, $\widetilde{w}_{it}$, capital intensity, $\widetilde{k_{it}/l_{it}}$, and labor productivity, $\widetilde{y_{it}/l_{it}}$, are then defined in logs as follows:

$$\widetilde{x}_{it} \equiv \log X_{it} - \overline{\log X_{-i,t}} \quad \text{where} \quad \overline{\log X_{-i,t}} \equiv \sum_{j \neq i} \frac{Y_{jt}}{\sum_{j \neq i} Y_{jt}} \log X_{jt} \quad \text{and} \quad X_{it} = W_{it}, \frac{K_{it}}{L_{it}}, \frac{Y_{it}}{L_{it}}. \quad (8)$$

While we omit the subscripts for industry and region for readability, it should be understood that the average wage and labor productivity for an establishment i's peers are defined within a given year, industry and state. These measures are centered around zero and denote an establishment's percentage deviation from the value added-weighted average of its peers. The advantage is that both relative measures are dimension-free metrics and can be compared across markets, years and industries.

Cross-section Our first exercise is to study the relationship between the labor share λ and its three components ($\widetilde{w}$, $\widetilde{y/l}$ and $\widetilde{k/l}$) in the cross-section. To do so, we run the following non-parametric regressions:

$$\widetilde{x}_{it} = f(\lambda_{it}) + \varepsilon_{it}, \qquad \widetilde{x}_{it} = \widetilde{w}_{it}, \ \widetilde{y_{it}/l_{it}}, \ \alpha\widetilde{k_{it}/l_{it}}, \qquad (9)$$

where $\widetilde{x}_{it}$ is either establishment i's relative wage, $\widetilde{w}_{it}$; labor productivity, $\widetilde{y_{it}/l_{it}}$; or (scaled) capital intensity, $\alpha\widetilde{k_{it}/l_{it}}$. The function $f(\cdot)$ is the object of interest: It indicates whether low-labor-share establishments pay lower wages than their peers, experience higher labor productivity and/or display higher capital intensity. To ensure that we measure economically-relevant relationships, each observation is weighted by the establishment's share in aggregate value added (the findings below are even stronger for unweighted regressions). Notice that we cannot include multiple right-hand-side variables in this local polynomial regression. Yet, since $\widetilde{w}$ and $\widetilde{y/l}$ are scaled within each industry, year and region, our results are unlikely to be driven by industry-, year- or region-specific effects.

The results of the three non-parametric regressions are displayed in the left panel of Figure 11. They paint two clear and striking pictures. First, relative wages are nearly orthogonal to the labor share: HP establishments do not on average pay their workers more or less than their peers.[10] By definition, differences in the labor share therefore have to be explained by differences in relative labor productivity. Indeed, the relationship between these two variables is strongly negative: $\widetilde{y/l}$

---

[9] We find that this definition of peer group strikes the right balance between making establishments comparable while keeping enough observations in a peer group to obtain sufficiently precise results. Choosing finer industry or region definitions do not change significantly the conclusions.

[10] Note that the error bands of our estimate denote the noise across establishments, not workers. Weighing observations (establishments) by their number of employees would reflect the more dispersed wage dispersion observed in worker-level or household level data. Even though we choose the more conservative establishment-level relative wage, the 95% error bands always include zero.

starts at about 1.6 for establishments with a near-zero labor share and then gradually declines through the labor share spectrum, hitting the average labor productivity $(\widetilde{y/l} = 0)$ at a labor share of $\lambda = 0.46$. These differences are large. For example, establishments with a labor share of $\lambda = 0.1$ experience a relative labor productivity of $\widetilde{y/l} = 1.04$. That means these establishments produce $\exp(1.04) \approx 2.83$ times more value added per worker than the average establishment in the same market, region and year. At the other end of the spectrum, establishments with a labor share of unity exhibit $\widetilde{y/l} = -0.61$, which means they produce only a bit above half the value added per worker $(\exp(-0.61) \approx 0.54)$ of the average establishment in the same market. HP establishments have an average relative labor productivity of 0.596 compared to -0.428 of Non-HP establishments; the average HP establishment thus produces about 2.8 times more value added per worker than the typical Non-HP establishment. Yet, in terms of relative wages, they do not differ at all.

Second, the nominal labor productivity profile is driven almost entirely by TFPR differences across establishments. While it is true that low-labor-share establishments tend to be on average more capital intensive than their peers, the contribution to labor productivity is small under our assumption of a constant-returns-to-scale Cobb-Douglas production function. From Equation (7), this implies that the remainder, nominal productivity, must be the main driver behind the large differences in labor productivity across the labor share spectrum.

Figure 11: Labor productivity dominates cross-sectional differences and time-series dynamics of labor shares of HP establishments.
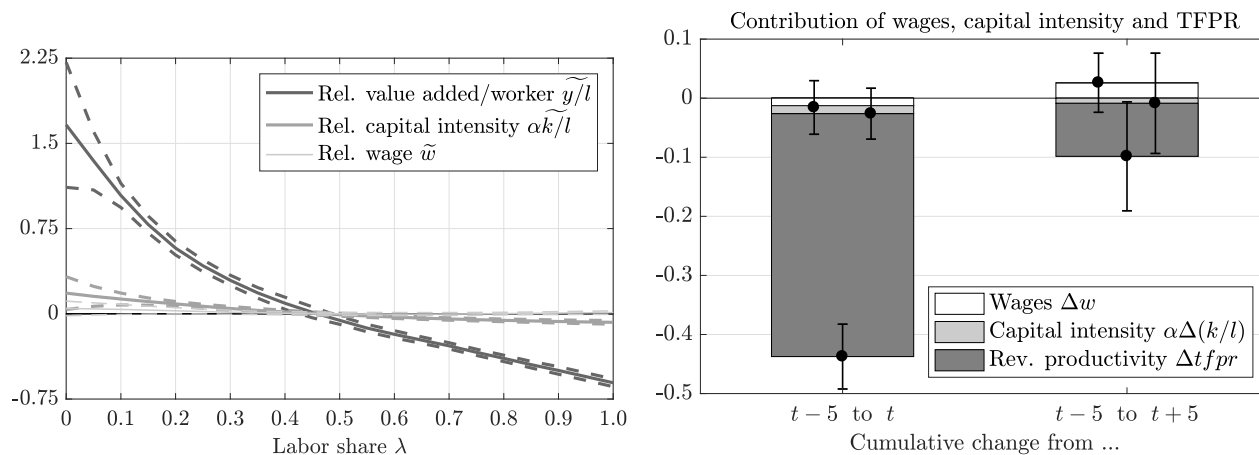
Note: Left panel displays the cross-sectional differences in relative value added per worker $\widetilde{y/l}$, the relative capital intensity $\alpha\widetilde{k/l}$ and the relative wage $\widetilde{w}$ against the labor share; we multiply the relative capital intensity by the typical capital elasticity in a constant-returns-to-scale Cobb-Douglas production function with $\alpha = 1/3$. All relative measures denote log-point differences vis-à-vis their peers as defined in Equation (8). Dashed lines denote 95% error bands.
Right panel displays the dynamic contributions of labor productivity growth and wage growth for labor share growth of the average HP establishment relative to their peers. The first bars display their cumulative contributions before $(t - 5$ to $t)$ and after $(t$ to $t + 5)$ the year an establishment is in HP status. Whiskers denote 95% error bands.

Dynamics   Next, we turn our attention to the dynamics of the labor share components. Hence, starting from the relationship in growth rates:

$$\log \lambda_{it}^{HP} = \log W_{it}^{HP} - \alpha \log(K_{it}^{HP}/L_{it}^{HP}) - \log TFPR_{it}^{HP}$$

we can apply the same regression strategy that was used to produce Figure 8 in order to decompose the change in the labor share of HP establishments at time t into the responses of wages, capital intensity and nominal productivity.[11] The right panel of Figure 11 shows that if TFPR was the main driver of cross-sectional differences along the labor share dimension, it also clearly is the main source of variation in the labor share of HP establishments. In the five years prior, the average HP establishment saw a labor share growth 43 ppt lower than that of the typical Non-HP establishment. This is almost entirely driven by a 42 ppt differential in TFPR growth, with no statistically significant action from wages and capital intensity. Incorporating the five following years, we see that the V-shape pattern of the labor share is a result of the retreat of the TFPR advantage of HP establishments following the initial jump. There is a slight relative wage gain, but it is not statistically different than zero. Unlike in the cross-section, variations in the capital intensity of HP establishments do not appear to play a significant role.

3.5   Finding 5: Employment responds less than before to positive TFPR shocks

In the previous section, we saw, among other things, that nominal labor productivity was central to understanding the labor share response of HP establishments. By definition, large fluctuations in labor productivity must imply that labor and value added do not move in lockstep. We now turn our attention to another striking fact: the responsiveness of employment to output has been markedly different during the recent period of declining aggregate labor share (2000s) relative to the early part of the sample when the labor share was more stable (1970s).

Before discussing employment dynamics, let us first contrast the cross-sectional relationship between labor share, wages and labor productivity in the 1970s and 2000s. We use the same non-parametric regression within-group approach from Equation (9), applied separately to each sample period. For the early part of the sample, labeled as "1970s," we focus on the 1967 through 1977 Censuses, while we use the 2002 through 2012 Censuses for the late part of the sample, denoted by "2000s."

The results are reported in the top row of Figure 12. A few things are worth noting. First, the link between wages and labor share was always weak, both economically and statistically. Second, while the inverse relationship between labor share and productivity was always strong, it became even more pronounced over time; the nonparametric estimate is particularly steep in the 2000s (top-right panel) for the establishments with the lowest labor share. Third, notice how the labor productivity curve crosses the x-axis at a lower labor share in the latter part of the sample. This

---

[11]Note that we study the components of the labor share *growth rate* $\Delta\lambda_{it} \equiv \log(\lambda_{it}/\lambda_{it-1})$, which is different from the absolute difference (in ppt) of the labor share, $d\lambda$, displayed in Figure 8. This explains why the total contribution of the three components in Figure 11 is not equal to the net numbers reported in Figure 8.

Figure 12: 1970s vs. 20002.

Rel. labor productivity and wages: 1970s

Rel. value added/worker $\widetilde{y/l}$
Rel. wage $\widetilde{w}$

Labor share λ

Rel. labor productivity and wages: 2000s

Rel. value added/worker $\widetilde{y/l}$
Rel. wage $\widetilde{w}$

Labor share λ

Contribution of labor, value added and wage growth: 1970s

Wages Δw
Labor Δl
Value added Δy

$t-5$ to t $t-5$ to $t+5$

Cumulative change from ...

Contribution of labor, value added and wage growth: 2000s

Wages Δw
Labor Δl
Value added Δy

$t-5$ to t $t-5$ to $t+5$

Cumulative change from ...

1970s

Net hiring

TFPR innovation

2000s

Net hiring

TFPR innovation

is another illustration of the massive reallocation we documented in Finding 1: The value-added-weighted average labor productivity takes place at a much lower labor share. Graphically, this corresponds to the dark grey line shifting downwards.

Next, we turn our attention to dynamics. We repeat the exercises of Equations (3) and (4), but this time applied separately to wages, employment and value added and study how HP establishments differ from their peers in the dynamics of these variables:

$$\Delta \log(\lambda_{it}^{HP}) = \Delta \log(W_{it}^{HP}) + \Delta \log(L_{it}^{HP}) - \Delta \log(Y_{it}^{HP}).$$

The middle panels of Figure 11 display the contribution of these three components to the labor share growth rate of HP establishments relative to non-HP establishments. The middle-left panel shows that in the 1970s, the majority of the adjustment in the five-year period preceding the current year was driven by a rise in value added (negative contribution to the labor share): the average HP establishment's labor productivity growth was 40 ppt higher than that of non-HP establishments. The relative change in labor share would have been more pronounced were it not for the fact that employment growth was 10 ppt higher for HP establishments. In the five following years, almost all the labor share growth differential disappears. This is mainly due to two factors: a retreat of value added following the time-t peak, but also a more robust relative response of employment for HP establishments whose hiring seems to respond to the strong value added growth, but with a delay. Ultimately, while the value added of HP establishments has clearly grown more over the 10-year span than that of their peers, the relative labor productivity is more or less back to where it was initially.

The dynamics in the 2000s are very different, at many levels. First, the value-added growth advantage of HP establishments between $t-5$ and t is larger, at 50 ppt instead of 40 ppt in the 1970s. Second, the V-shaped pattern is now more pronounced: not only is the value added growth differential sharper initially, but it now is only 17 ppt after 10 years, compared to 22 ppt in the 1970s. Third, the response of employment is strikingly different from the early part of the sample: Between $t-5$ and t, employment growth is 7 ppt *lower* for HP establishments relative to their Non-HP peers, despite the sharp increase in value added. By $t+5$, the cumulative employment growth differential is indistinguishable from zero.

Such a disconnect between value added and labor input, particularly in the latter part of the sample, is surprising: standard models would predict that high-productivity establishments expand their workforce or at least increase their wages. Yet, it is in line with recent work documenting the decline in the responsiveness of the economy to shocks, see for example Decker, Haltiwanger, Jarmin, and Miranda (2017a,b); Cooper, Haltiwanger, and Willis (2017); Ilut, Kehrig, and Schneider (2014).

To further analyze this disconnect, we follow the empirical setup in Ilut, Kehrig, and Schneider (2014) and non-parametrically estimate net hiring as a function of TFPR shocks

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

where z_{it} is the TFPR 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, using annual data from the Annual Survey of Manufactures (ASM). We aggregate the estimated hiring responses to obtain the typical hiring response of heterogeneous establishments within an industry and decade.

The non-parametric estimates shown in the bottom panels of Figure 12 clearly indicate that the employment-TFPR relationship has changed markedly over this forty-year period. While hiring in the 1970s is almost linear in TFPR shocks, it becomes highly asymmetric by the 2000s. This asymmetry is entirely driven by a weaker response of employment to positive shocks[12]: by the 1990s, the response to a two-standard deviation TFPR shock ($z \approx +0.4$) becomes statistically insignificant at the 5% level (bottom panels). Negative shocks, in contrast, lead to significant firing throughout the sample.

It should be noted that this asymmetry is compounded by the fact that at the same time, the distribution of TFPR shocks changes. As Kehrig (2011) has shown (see Fig. 12 in that paper), the long-run cross-sectional dispersion of productivity levels becomes wider with every recession since the 1980/82 recession. This means there are more firms in the tails of the TFPR innovation distribution in Figure 12. This increasing polarization of productivity shocks implies that the asymmetry bears out more strongly in the 2000s than in the 1980s.

Why don't establishments that become highly productive hire more? Why have they become less responsive over time? The transient nature of the labor share and value added may provide an answer: 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 positive realization of TFPR that it expects to be temporary.

3.6 Finding 6: Low labor shares are related to a "product price premium"

The previous sections revealed a number of puzzling findings. Not only did we show that the aggregate labor share declined due to a dramatic reallocation of value added towards "hyper-productive" establishments, we also discovered low labor shares at the micro level to be temporary phenomena. Yet, a fundamental question remains: what drives the labor share dynamics of these "mayfly HP establishments"? Finding 4 gave us a hint about the cause: cross-sectional and dynamic differences between HP and Non-HP appear to be driven by *nominal* value added per worker, particularly TFPR. This leaves two candidate forces driving the aggregate labor share decline: nominal price dynamics and real technological change. In this section, we attempt to disentangle these two opposite forces and show that demand-side factors appear to be a key driver of micro-level labor share patterns.

---

[12]Gutiérrez and Philippon (2017) document a similar pattern for investment.

3.6.1 Measuring prices

The previous sections showed that HP status was mostly a temporary phenomenon: these establishments experienced a marked decline in their labor share over the previous five years, yet their labor share dynamics do not look significantly different than that of their peers over a longer period. We found that these dynamics were driven mainly by sharp fluctuations in labor productivity, and in particular TFPR. These two concepts are *nominal* variables, however. This implies that they may be driven by either prices or physical productivity.

In order to identify the relative contributions of these two distinct forces, we turn to another data source provided the U.S. Census Bureau: the product trailer to the Census of Manufactures. For each establishment, the product trailer records the value of sales generated by individual products (variable PV). These sales numbers are available for the vast majority of establishments and allow us to study how much of an establishment's sales stems from a few "superstar products" or from the broad "brand appeal" that lifts the sales of all products equally. In addition, the product trailer collects information on the physical quantity of products shipped (variable PQS) for a sample of establishments, whenever a meaningful metric can be used. In those cases, we can compute the average product-level price charged by an individual establishment. We will use this subset of the database to disentangle the contribution of prices from that of physical productivity.

Our analysis is inspired by the approach pioneered in Foster, Haltiwanger, and Syverson (2008), though we deviate from their methodology in that we consider products at the 10-digit NAICS level, a finer definition of product than most of the literature.[13] This is a product-coding system devised by Census and based on the NAICS industry code. Second, because our aim is to study an establishment's prices and real productivities relative to that of its peers, we only use observations that are not imputed to ensure that values are directly comparable (for details see Appendix A.4).[14]

The price data have some drawbacks, however. For one, the imputation flags for prices and quantities are only available starting with the 1992 Census, and coverage is very limited in the 1992 and 2012 Censuses. Most importantly, only a few industries have well-defined quantity measures for (a subset of) their products. In addition to the products studied by Foster, Haltiwanger, and Syverson (2008), examples of the products we consider are certain homogeneous chemicals (measured in metric tons) or metals such as aluminum sheets (measured in thousand lbs), for example, but not vehicles or clothing which are measured in the generic unit "number." All these limitations imply that we are left with a panel of 130 thousand year-establishment-product observations whose quality is high enough to study separately prices and quantities. We call the resulting panel the "Matched Price Sample" to distinguish it from the "Full Census Sample," our default panel.

The Matched Price Sample allows us to link an establishment's product-level prices and its revenue labor productivity, which we found to be the key driver of labor shares in the cross section

---

[13]Foster, Haltiwanger, and Syverson (2008) define products at the 7-digit SIC code level while Bernard, Redding, and Schott (2010, 2011) aggregate product sales to the 5-digit SIC level of products; both definitions are coarser than ours.

[14]White, Reiter, and Petrin (2018) have shown that the product trailer dataset is seriously contaminated by imputations based on industry averages or regression models.

and time series. Since all price data are sales based, we are switching to studying sales per worker rather than value added per worker when analyzing the price-vs-physical-productivity difference. We define relative sales per worker analogous to that of relative value added per worker in Equation (8):

$$\widetilde{p_{it}q_{it}/l_{it}} \equiv \log\left(P_{it}Q_{it}/L_{it}\right) - \overline{\log(P_{-i,t}Q_{-i,t}/L_{-i,t})} \tag{10}$$

$$\text{where } \overline{\log(P_{-i,t}Q_{-i,t}/L_{-i,t})} \equiv \sum_{j\neq i}\frac{P_{jt}Q_{jt}}{\sum_{j\neq i}P_{jt}Q_{jt}}\log\left(P_{jt}Q_{jt}/L_{jt}\right).$$

Naturally, the establishments in the Matched Price Sample are more homogeneous in the type of products than those in the Full Census Sample. The distribution of $\widetilde{p_{it}q_{it}/l_{it}}$ can thus be expected to be more compressed in the Matched Price Sample than in the Full Census Sample. Yet, our analysis in Appendix A.5 reveals that differences in sales per worker remain the main driver of both cross-sectional and dynamic moments of the labor shares in the Matched Price Sample.

3.6.2 Product prices across establishments and over time

In order to make prices comparable across establishments, we follow the treatment of nominal wages and labor productivity in Section 3.4 by comparing establishment-level prices to a peer group. This time, however, we have to start at the product level. First, we normalize prices at the level of the 10-digit NAICS product ℓ:

$$\widetilde{p}_{i\ell t} \equiv \log P_{i\ell t} - \overline{\log P_{-i,\ell t}} \qquad \text{where } \overline{\log P_{-i,\ell t}} \equiv \sum_{j\neq i}\frac{P_{j\ell t}Q_{j\ell t}}{\sum_{j\neq i}P_{j\ell t}Q_{j\ell t}}\log P_{j\ell t} \tag{11}$$

That is, we compare the price of product ℓ sold by establishment i at time t to the weighted average of the prices charged for the same product by all other establishments $j \neq i$ in the same year. $\widetilde{p}_{i\ell t}$ denotes the log-point difference that establishment i charges for product ℓ compared to the average price charged by its peers for the same product.

Next, we aggregate these relative prices across all products offered by establishment i and year t to obtain the establishment-level sales-weighted average relative product price $\widetilde{p}_{it}$:

$$\widetilde{p}_{it} \equiv \sum_{\ell \in i}\widetilde{p}_{i\ell t}\frac{P_{i\ell t}Q_{i\ell t}}{\sum_{\ell \in i}P_{i\ell t}Q_{i\ell t}}.$$

We refer to $\widetilde{p}_{it}$ as the average "product price premium" that establishment i charges relative to its peers across all the products it sells. This measure represents the average log-point difference between an establishment's output prices and those of its peers.[15]

---

[15]A word of caution is warranted here: as argued by Edmond, Midrigan, and Xu (2018), the theoretically correct approach would be to use a cost-weighted average. In our case, unfortunately, the lack of cost information at the product level means that we have no choice but to rely on a sales-weighted average.

Similar to our earlier approach, we non-parametrically estimate the cross-sectional relationship between the product price premium and the labor share. Because sales are multiplicative in prices and quantities, we can interpret the magnitude of the product price premium as the share of relative sales per worker explained by prices; the remainder is the portion explained by physical labor productivity $\widetilde{q/l}$. The contributions of these two components to differences in relative sales are depicted in the left panel of Figure 13.

In addition, we run regressions analogue to Equations (3) and (4) for the relative price $\widetilde{p}$ and the relative productivity $\widetilde{pq/l}$ measures. The cumulative dynamics of relative prices and relative physical productivities derived from these regressions are displayed in the right panel of Figure 13.

Figure 13: Relative labor productivity and relative prices

Note: Left panel displays the cross-sectional differences in relative prices $\widetilde{p}$ (dark grey bars) and relative physical labor productivity $\widetilde{q/l}$ (light grey bars) against the labor share; $\widetilde{p}$ defined in Equation (11) and $\widetilde{q/l}$ is defined as the ratio of $\widetilde{pq/l}$ (defined in Equation (10)) and $\widetilde{p}$.
Right panel displays the dynamic contributions of the growth in relative prices $\Delta\widetilde{p}$ and labor productivity growth $\Delta(q/l)$ for sales per worker growth $\widetilde{pq/l}$ of the average HP establishment relative to their peers. The first bars display their cumulative contributions before ($t-5$ to t) and after (t to $t+5$) the year an establishment is in HP status. Whiskers denote 95% error bands.

We can see from the left panel that relative sales per worker are driven by both price and physical labor productivity differences, but the role of prices is crucial in characterizing those establishments with the lowest labor share. For example, for establishments with a labor share below 20%, relative prices explain more than two thirds of the sales per worker differences $(\widetilde{pq/l})$. However, relative prices play only a little role in explaining differences in sales/worker of establishments with below-average sales per worker – those with a labor share of 50% and more. The average relative price of HP establishments is 0.15 compared to -0.041 for Non-HP establishments. This difference implies a product price premium for HP establishments of roughly $\exp(0.15 + 0.041) \approx 21\%$. This contributes a fair amount to the relative sales per worker that HP establishments generate to their peers of $\exp(0.430 + 0.096) \approx 69\%$.

The right panel of Figure 13 displays the dynamic analysis. Analogously to the regression in Equations (3) and (4), we estimate how much HP establishments differ in their *dynamics* of

29

price and sales per worker from Non-HP establishments. Again, we infer the physical productivity growth of HP establishments from the difference between the two. The results of the dynamic analysis are even starker than those from the cross-section: compared to their Non-HP peers, relative prices of HP establishments increase by 16.8% from the previous Census year (from $t - 5$ to t). In the subsequent five years most, but not all, of that jump in the product price premium is reverted: the change in average product price premium from $t - 5$ to $t + 5$ is 7.8% higher for those establishments that are HP at time t relative to their peers. While the HP-vs-Non-HP price dynamics are significantly different from zero, this is not the case for physical productivity: the relative cumulative growth rates of $\widetilde{q/l}$ hover around -2% and are statistically insignificant.

All in all, the evidence in this section indicates that the term HP we coined for establishments with very low labor shares in fact refers to Hyper-*Revenue*-Productive establishments[16]: their status seems to be mainly derived from a capacity to extract a lot of revenue out of their workforce.

How much do these price dynamics contribute to the labor share dynamics of HP establishments? We can isolate the contribution of price changes to the growth rate of the labor share as follows:

$$\Delta \log \lambda_{it}^{HP} = \Delta \log w_{it}^{HP} - \Delta \log \left(\frac{y_{it}^{HP}}{l_{it}^{HP}} \right)$$
$$= -\Delta \log p_{it}^{HP} \frac{p_{it}^{HP} q_{it}^{HP}}{y_{it}^{HP}} + Residual_{it}$$

where $Residuals_{it}$ collects the remaining contributions of wage growth, physical output and intermediate inputs per worker. In the Matched Price Sample, we find that HP establishments experienced a price growth of 14.9% on average over the past five years, while their labor share contracts by 52.6% over the same horizon. In addition, the ratio of sales to value added, $\frac{p_{it}^{HP} q_{it}^{HP}}{y_{it}^{HP}}$, equals about 2.54 for HP establishments. This means that a typical price increase implies a $14.9 \times 2.54 \approx 37.8$ ppt contribution, which is roughly three quarters of the total labor share decline of HP establishments. While we can carry out this calculation only in the Matched Price Sample, prices in the Full Census Sample are likely more heterogeneous and may account for an even larger share.

4 Counterfactuals

The evidence from the previous sections has highlighted a number of striking cross-sectional and dynamic features of the micro-level data. In this section, we wish to explore the potential roles of two of these features in explaining the sharp decline of the aggregate labor share: the evolution of the V-shaped pattern over time (Finding 2), the fact that low labor shares have become increasingly correlated with past establishment size (Finding 3), and the declining responsiveness of employment to positive value added shocks (Finding 5). We start by explaining our framework before studying

---

[16]In Appendix B.2 we study the potential role of transfer prices across establishments within firms, but find that labor share patterns at the firm and establishment levels resemble each other.

a number of scenarios.

Our first approach relies on the HP-vs-Non-HP distinction that we introduced earlier, with a particular emphasis on transitions. Specifically, in a given period t, we define four types of establishments:

Type 1: Establishments that were Non-HP at $t-1$ and remain Non-HP at t

Type 2: Establishments that were Non-HP at $t-1$ and become HP at t

Type 3: Establishments that were HP at $t-1$ and become Non-HP at t

Type 4: Establishments that were HP at $t-1$ and remain HP at t

For a given group j in period t, we denote its current and lagged value-added weights and labor shares as: ω_t^j, $\omega_{t,lag}^j$, λ_t^j and $\lambda_{t,lag}^j$. As we explain next, we can extract from these values the (relative) shocks to value added and the wage bill. These will form the reference points for our counterfactuals. Note that throughout this section, we focus on the same balanced panel of establishments that was used to produce Figure 5. This is because we need to maintain longitudinal consistency from Census to Census and – since we concatenate our dynamic counterfactuals – throughout the entire sample. Of the about 64 thousand establishment-year observations, about 70% are Type 1 establishments, the other three types account account for roughly equal shares of the remainder.

4.1 Extracting Shocks from the Data

Let us denote by $\widehat{\alpha}_t^j$ and $\widehat{\alpha}_t$ the shocks to group-j's and aggregate value added, respectively (all variables with values derived from the data are hatted). This allows us to rewrite the value-added (VA) weight of group j at time t as:

$$\widehat{\omega}_t^j = \frac{\widehat{VA}_t^j}{\widehat{VA}_t} = \frac{\widehat{\alpha}_t^j \widehat{VA}_{t,lag}^j}{\widehat{\alpha}_t \widehat{VA}_{t,lag}} = \frac{\widehat{\alpha}_t^j}{\widehat{\alpha}_t} \widehat{\omega}_{t,lag}^j \tag{12}$$

which implies that the VA shock for group j can be computed as:

$$\widehat{\alpha}_t^j = \widehat{\alpha}_t \frac{\widehat{\omega}_t^j}{\widehat{\omega}_{t,lag}^j} \tag{13}$$

Since all that will matter for our counterfactuals are relative changes across types, we assume for simplicity that $\widehat{\alpha}_t = 1$. This allows us to directly recover the (relative) VA shocks $\widehat{\alpha}_t^j$ from the ratios of weights.

Next, let us now write the ratio of the current and lagged labor shares of group j as:

$$\frac{\widehat{\lambda}_t^j}{\widehat{\lambda}_{t,lag}^j} = \frac{\widehat{WB}_t^j}{\widehat{WB}_{t,lag}^j} \frac{\widehat{VA}_{t,lag}^j}{\widehat{VA}_t^j} = \frac{\widehat{WB}_t^j}{\widehat{WB}_{t,lag}^j} \frac{1}{\widehat{\alpha}_t^j} = \frac{\widehat{\gamma}_t^j}{\widehat{\alpha}_t^j}$$

where $\widehat{\gamma}_t^j$ is the wage-bill shock of group j. This shock can be recovered by using the known values $\widehat{\alpha}_t^j$, $\widehat{\lambda}_t^j$ and $\widehat{\lambda}_{t,lag}^j$.

Finally, our approach requires that we recover the relative weights of the two types coming from the same pool last period (HP or Non-HP). Specifically, we denote the relative value-added of time-t Non-HP establishments (type 1) within the pool of $t-5$ Non-HP establishments by $\widehat{\xi}_t^{NHP}$; and by $\widehat{\xi}_t^{HP}$ the relative value-added weight of time-t Non-HP establishments (type 3) within the pool of $t-5$ HP establishments. The empirical values can be recovered directly from the weights:

$$\widehat{\xi}_t^{NHP} = \frac{\widehat{\omega}_{t,lag}^1}{\widehat{\omega}_{t,lag}^1 + \widehat{\omega}_{t,lag}^2} \tag{14}$$

$$\widehat{\xi}_t^{HP} = \frac{\widehat{\omega}_{t,lag}^3}{\widehat{\omega}_{t,lag}^3 + \widehat{\omega}_{t,lag}^4} \tag{15}$$

The objects ξ_t^{NHP} and ξ_t^{NHP} are directly related to the correlations between current labor share and past size that we documented in Section 3.3. A declining value of $\widehat{\xi}_t^{NHP}$, for example, would imply that of the total value added accounted for by establishments that were Non-HP at $t-5$, a large (small) portion belongs to establishments that became HP (remained Non-HP) at t. In other words, establishments that acquired HP status tended to be larger over time. In line with the earlier evidence, this is indeed what we observe in this balanced panel: out of the establishments that were Non-HP in 1967, 5.6% in value-added share ended up as HP in 1972 (that is, $\widehat{\xi}_{1972}^{NHP} = 0.944$). By 2012, this value had more than tripled to 18.3% ($\widehat{\xi}_{2012}^{NHP} = 0.817$).

4.2 Building the Counterfactuals

We now turn our attention to constructing counterfactual scenarios, with the objective of exploring the respective roles of shocks (α_t^j, γ_t^j) and selection (ξ_t^{NHP}, ξ_t^{HP}). Our approach is sequential, starting from 1972 and taking the 1967 values as given. The main steps are described below and illustrated graphically in Figure 14.

1. Construct counterfactual time series for the shock (α_t^j, γ_t^j) and selection (ξ_t^{NHP}, ξ_t^{HP}) variables. In some instances, these will simply be kept equal to the empirical time series described above.

2. Compute the 1972 weights using the counterfactual VA shocks.

$$\omega_{1972}^j = \frac{\alpha_{1972}^j \omega_{1972,lag}^j}{\sum_{k=1}^4 \alpha_{1972}^k \omega_{1972,lag}^k}$$

3. Compute the 1972 labor shares using the counterfactual VA and wage bill shocks.

$$\lambda_{1972}^j = \frac{\lambda_{1972,lag}^j \gamma_{1972}^k}{\alpha_{1972}^k}$$

32

4. Compute the counterfactual aggregate labor share for 1972.

$$\lambda_{1972} = \sum_{j=1}^{4} \omega_{1972}^{j} \lambda_{1972}^{j}$$

5. Compute the lagged weights and lagged labor shares to be used in the next period (1977).

    (a) Compute the aggregate weights for Non-HP and HP plants in 1972

    $$\omega_{1972}^{NHP} = \omega_{1972}^{1} + \omega_{1972}^{3} \qquad \omega_{1972}^{HP} = \omega_{1972}^{2} + \omega_{1972}^{4}$$

    (b) Determine the lagged weight of each type for 1977, using the selection variables:

    $$\omega_{1977,lag}^{1} = \xi_{t}^{NHP} \omega_{1972}^{NHP} \qquad \omega_{1977,lag}^{2} = (1 - \xi_{t}^{NHP}) \omega_{1972}^{NHP}$$

    $$\omega_{1977,lag}^{3} = \xi_{t}^{NHP} \omega_{1972}^{HP} \qquad \omega_{1977,lag}^{4} = (1 - \xi_{t}^{HP}) \omega_{1972}^{NHP}$$

    (c) Compute the lagged labor share of each type for 1977.[17]

6. Repeat steps 2 to 5 for the other Census years.

Next, we apply this methodology to compare counterfactual and actual aggregate labor shares under various scenarios.

4.3 Results

We focus on three types of counterfactuals. For the first set, we modify the time series of shocks to value added and the wage bill (α_{t}^{j}, γ_{t}^{j}). In the second, we adjust their relative importance to account for changes in responsiveness. Finally, in the third set, the selection into Non-HP and HP status (ξ_{t}^{NHP}, ξ_{t}^{HP}) is altered. The counterfactual aggregate labor shares for all scenarios are shown in Figure 15, against the actual labor share.

4.3.1 Shock process

In the first exercise, we simply turn off all shocks ($\alpha_{t}^{j} = \gamma_{t}^{j} = 0$ for all j and $t = 1972, ..., 2012$), keeping selection unchanged ($\xi_{t}^{NHP} = \widehat{\xi}_{t}^{NHP}$, $\xi_{t}^{HP} = \widehat{\xi}_{t}^{HP}$). The counterfactual labor share is shown in the top panel. Not surprisingly, the aggregate labor share remains constant in this case.

---

[17]Throughout our exercises, we assume that the ratios of the labor shares of the two subgroups (1 vs 2, and 3 vs 4) remain the same as in the original data. Unlike other variables, we find no clear trend in these ratios over time. This assumption implies that for Type 1, for example, the expression is:

$$\lambda_{1977,lag}^{1} = \lambda_{1972}^{NHP} \frac{\omega_{1977,lag}^{1} + \omega_{1977,lag}^{2}}{\omega_{1977,lag}^{1} + \omega_{1977,lag}^{2} \frac{\widehat{\lambda}_{1977,lag}^{2}}{\widehat{\lambda}_{1977,lag}^{1}}}$$

Figure 14: Methodology for constructing the counterfactuals
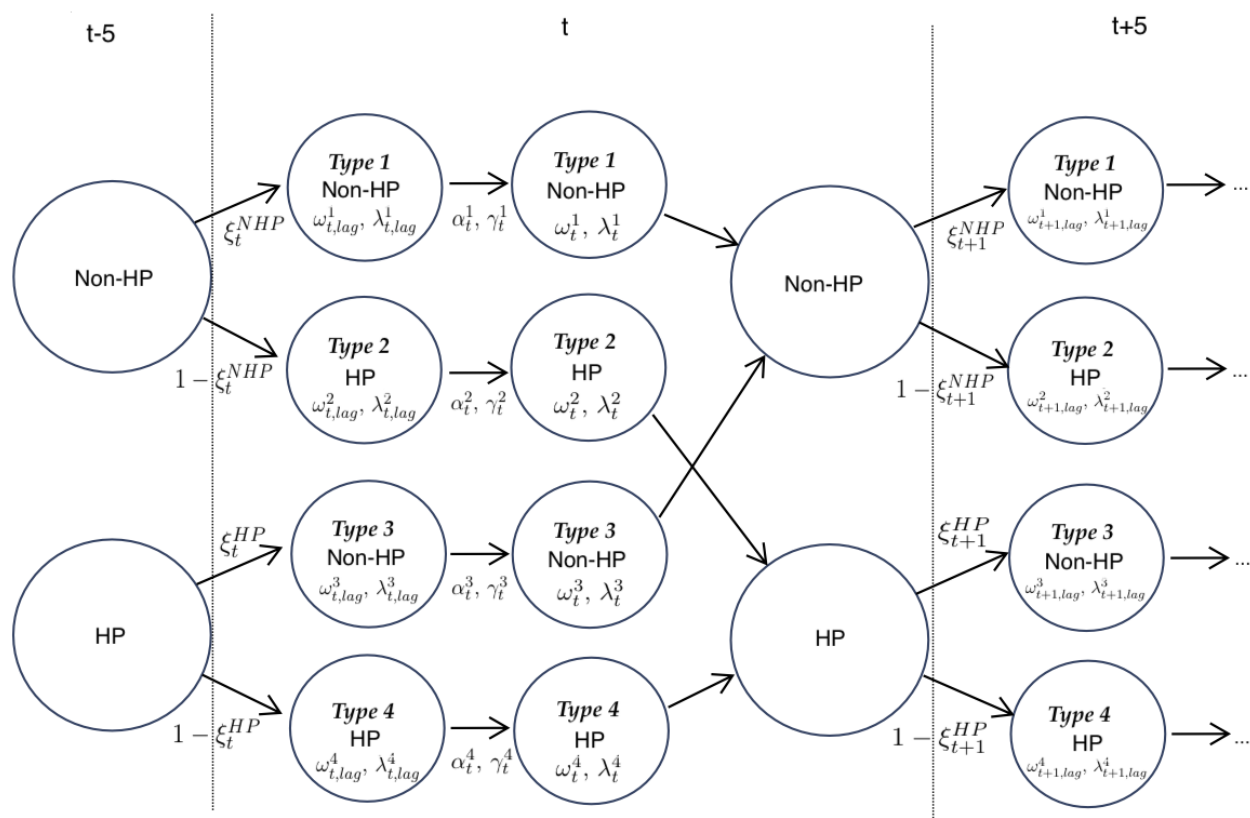
Recall that in Section 3.2.3 we documented a deepening over time of the labor share V-shapes for HP establishments: the absolute changes in the labor share in and out of HP status have become larger. In the next exercise, we freeze both the VA and wage bill shocks of types 2 (Non-HP to HP) and 3 (HP to Non-HP) to their 1972 values, as well as impose symmetry of the V-shape.[18] The results are found in the middle-left panel of Figure 15: while this time the decline in the counterfactual aggregate labor share is smaller than the actual, the impact is quantitatively limited (13 ppt vs 15 ppt). This is consistent with the finding of Decker, Haltiwanger, Jarmin, and Miranda (2017b) who find a limited role for changes in the shock process to explain the decline in aggregate labor productivity, an object intimately linked to the labor share.

4.3.2 Responsiveness to shocks

We documented in Section 3.5 that the response of employment to positive TFPR shocks had become much more subdued over our sample period. Responsiveness to negative shocks, on the other hand, appears not to have changed significantly. As we noted earlier, some researchers have attributed recent macroeconomic phenomena to the decline in firm responsiveness to shocks.

In the context of our counterfactual setup, establishments of type 2 are those that are hit by large positive shocks, making them move from Non-HP to HP status. They should therefore be the ones predominantly affected by the decline in employment responsiveness that we documented earlier. For the next exercise, we therefore simulate the aggregate labor share under the assumption that the shocks buffeting type-2 establishments retain the same proportionality as in 1972. More specifically, this counterfactual assumes that wages bills remain as responsive to value added shocks as they used to in the 1970s: we keep the VA shocks to their actual values ($\alpha_t^2 = \alpha_{1972}^2$) but choose the wage bill shocks γ_t^2 such that $\alpha_t^2/\gamma_t^2 = \alpha_{1972}^2/\gamma_{1972}^2$.

The result of this exercise is shown in the middle-right panel of Figure 15. While the counterfactual aggregate labor share follows initially very closely the actual series, a significant divergence starts appearing in the early 2000s. By 2012, the labor share is higher than in the data by about 6 ppt (0.48 vs. 0.42), out of a total actual decline of 15 ppt (0.57 to 0.42). Declining responsiveness therefore appears to be a potential factor behind the decline in the aggregate labor share, at least for the last third of our sample.

4.3.3 Selection

Next, we turn our attention to the role of selection. As we discussed in Section 3.3, the likelihood of being selected as a HP establishment from one period to the next has become increasingly tilted towards larger establishments.

In the next counterfactual exercise, we force the selection from Non-HP into HP to be constant over time, equal to its 1972 value. In other words, we assume that $\xi_t^{NHP} = \hat{\xi}_{1972}^{NHP}$ for all $t =$

---

[18]Specifically, we assume (1) that $\alpha_t^j = \alpha_{1972}^j$ and $\gamma_t^j = \gamma_{1972}^j$ for $j = 2, 3$ and $t = 1977, ..., 2012$; and (2) that $\alpha_t^3 = 1/\alpha_t^2$ and $\gamma_t^3 = 1/\gamma_t^2$ from 1977.

Figure 15: Counterfactual exercises

No VA nor wage shocks

Constant and symmetric V-shape

Constant responsiveness to positive VA shocks

Constant selection out of NHP

Constant selection out of both NHP and HP
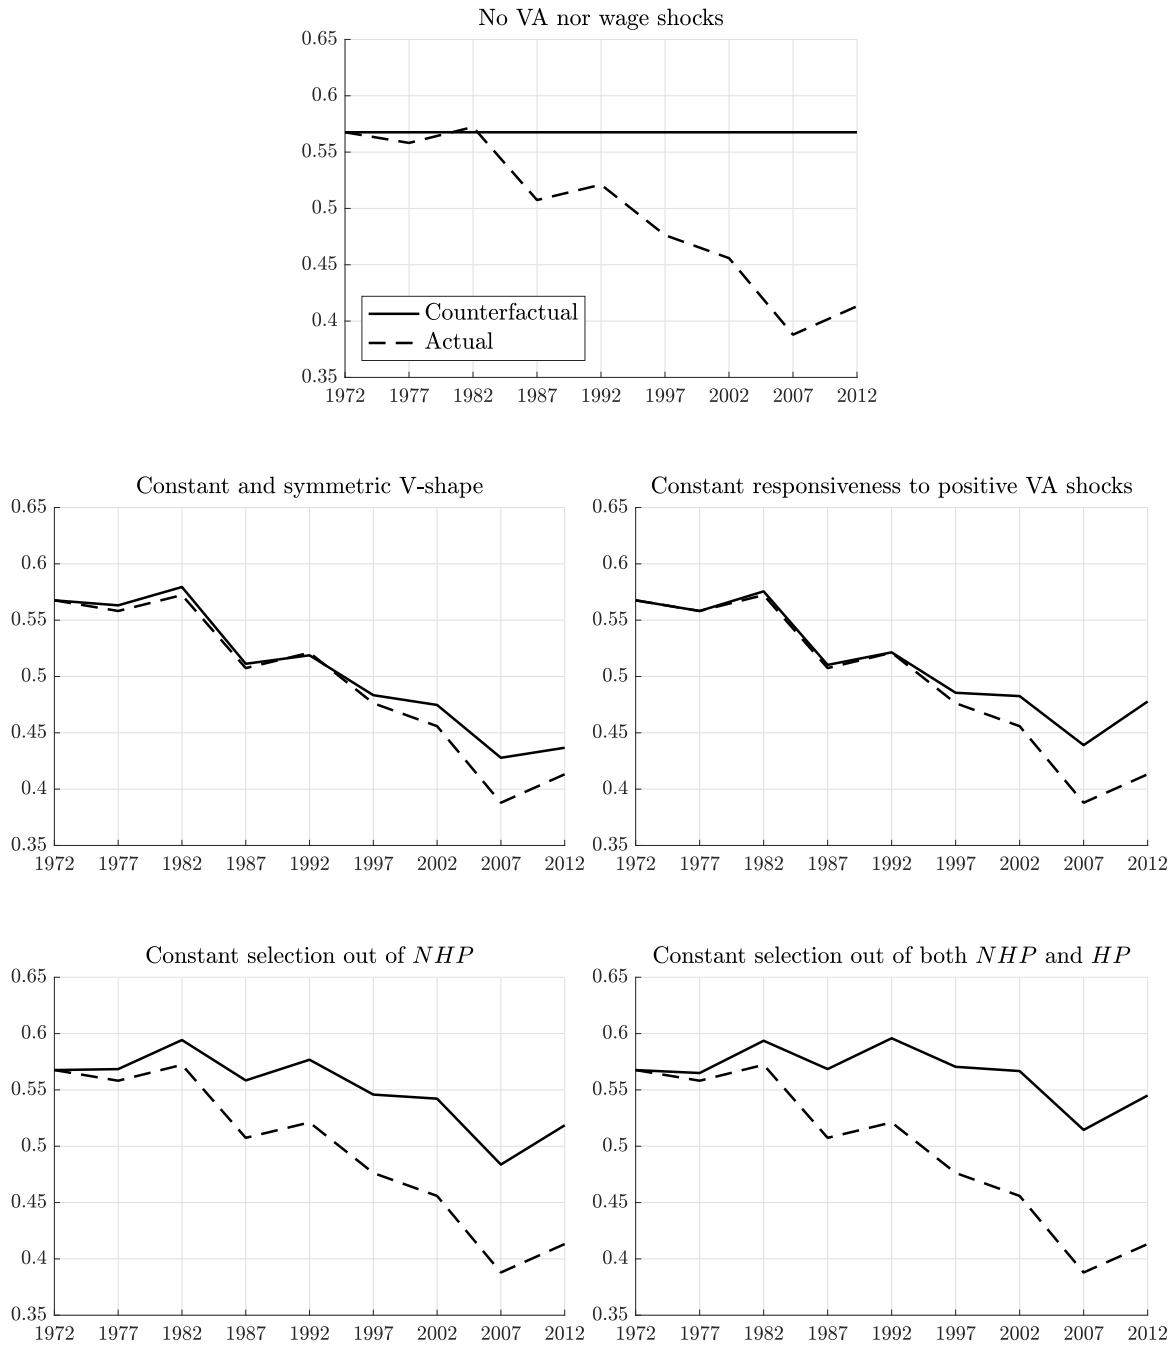
Note: This figure displays counterfactual exercises that assess the effect of various shocks: No changes to both wage bills and value added would have kept the labor share constant (top panel), while freezing the responsiveness of the wage bill to the positive value-added shock for type 2 (Non-HP to HP) would have reduced by 6 ppt the decline in the labor share (middle-right panel). Keeping the V-shape pattern of HP establishments' labor share would have little impact on the aggregate (middle-left panel). Keeping the selection out of Non-HP (bottom left) or both non-HP and HP (bottom right) establishments would have eliminated the cumulative decline in the aggregate labor share by about 65% and 85% respectively.

1977, ..., 2012. The impact of this margin on the aggregate labor share is striking: in the bottom-left panel of Figure 15, we observe that the decline under this counterfactual scenario is only about 5 percentage points (from 0.57 to 0.52), compared to an actual drop of 15 ppt (0.57 to 0.42). The role of selection appears to become quantitatively relevant as early as the mid-1980s, and the effects get compounded throughout the 1990s and 2000s.

Finally, we repeat the exercice by also freezing the selection out of HP ($\xi_t^{HP} = \widehat{\xi}_{1972}^{HP}$). The picture in the bottom-right panel is similar to the one in the bottom-left: while the drop in the aggregate labor share is somewhat more muted (from 0.57 to 0.55), most of the action clearly comes from the increasingly tilted selection into HP status towards establishments that account for a larger share of value added.

In sum, our counterfactual exercises appear to indicate that the role of value added and wage bill shocks in explaining the decline in the aggregate labor share is limited, at most explaining one third of the change. On the other hand, we found that the change over time in the correlation between the size of a establishment and its future labor share can explain up to 85% of the fall in the labor share between 1972 and 2012.

5   Potential forces behind the labor share decline

Our micro-level anatomy of the labor share in Section 3 has unearthed a number of striking facts that should help researchers in devising mechanisms and models that can structurally account for this aggregate phenomenon. In what follows, we investigate a few more potential factors and forces to guide us further in that direction.[19]

5.1   Components of labor compensation

The labor cost variable used in the numerator of the labor share contains various components. In the Census data, it is possible to distinguish between production worker wages, salaries for non-production workers as well as ancillary labor costs. A natural theory of the labor share decline could be skill-biased technical change which likely would disproportionately hurt a particular type of labor. If robots and production labor were substitutes, then one would expect capital-embodied technical change reduce the portion of labor compensation going to production labor. Skilled workers are likely more complementary to capital, so their salaries should not be as affected.

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

---

[19]Disclosure limitations restrict the results in this section to the 1967-2007 period.

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 whether 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 that are shown in Figure 16 and Table 3:

$$\lambda_t = \underbrace{\frac{w_t^{pw} L_t^{pw}}{Y_t}}_{\text{Wage bill}} + \underbrace{\frac{w_t^{npw} L_t^{npw}}{Y_t}}_{\text{Salaries}} + \underbrace{\frac{w_t^{ben} L_t^{ben}}{Y_t}}_{\text{Ancill. labor costs}} . \tag{16}$$

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

Figure 16: Dynamics of labor share components

Note: This figure displays three portions of aggregate labor compensation: wages of production workers, salaries of non-production workers and total ancillary labor costs such as unemployment insurance and health benefits. The secular decline of the production worker wage bill was first compensated by a rise in the ancillary labor compensation until the early 1980s when all three portions start to decline.

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

Component	1967-2007	1967-1982	1982-2007
	(percentage point changes)		
Aggregate 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

Note: Results from the shift-share decompositions as defined in (17) applied to the three types of labor compensation listed in Equation (16). The acceleration of the labor share decline almost exclusively stems from a more negative within-group adjustment term in salaries and ancillary labor costs suggesting that all types of labor suffer.

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. Beyond that point, 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.

5.2 The decline in the labor share occurs within, not between, industries, regions and types of firms

We now turn our attention towards factors that we expect would generate an important reallocation of economic activity across sectors, regions or types of establishments. To keep our conclusions broad, our exercises are built around a general within/between decomposition approach.

5.2.1 Industry factors

The observed fall in the aggregate labor share could stem from a subset of specific industries that experience a decline while others exhibit only a small change in their labor share. Or, 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 within- and between-industry components using Equation (17):

$$\Delta\lambda_t = \underbrace{\sum_j \Delta\lambda_{jt}\omega_{jt-1}}_{\text{Within adjustment}} + \underbrace{\sum_j \lambda_{jt-1}\Delta\omega_{jt}}_{\text{Between reallocation}} + \underbrace{\sum_j \Delta\lambda_{jt}\Delta\omega_{jt}}_{\text{Residual}} \tag{17}$$

where λ_j denotes the industry-level labor share and ω_j the share of value added accounted for by industry j.

Panel A. in Table 4 displays the results from this industry-level decomposition. It shows that most of the labor share decline between 1967 and 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. Importantly, the acceleration of the labor share decline starting in the 1980s is predominantly captured by the within-industry adjustment term, with a much more limited role for between-industry reallocation. Considering instead 4-digit NAICS industries (not displayed) does not change this takeaway.

Table 4: Labor share declines within and between industries, regions, legal forms of organization

Portions of labor share change	1967-2007	1967-1982	1982-2007
	(percentage point changes)		
Aggregate labor share change	−4.9	−0.9	−7.3
A. 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
B. Census regional 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
C. Legal form of organization			
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
D. Public vs. private firms			
Within-group adjustment	−5.1	−0.5	−7.9
Between-group reallocation	+0.2	−0.5	+0.5
Residual	+0.1	+0.0	+0.1

Note: Results from the shift-share decompositions as defined in (17) applied to indusitres (Panle A.), regions (Panel B.), legal forms of organizations (Panel C.) and the set of publicly traded versus privately held firms (Panel D.). The acceleration of the labor share decline almost exclusively stems from a more negative within-group adjustment term suggesting that reallocation between these groups only plays a minor role.

5.2.2 Regional factors

As in the within-between industry decomposition of the labor share change, one could study regional factors. This channel should be predominant if firms sort into different regions according to their labor share. There is a vast array of reasons how some regions could have a different effect

on the labor share than others. For example, states may provide tax incentives if firms open a new establishment in their state. If labor laws in these states are friendlier to businesses, then workers may not be compensated as much as they are in other states. With the aim of determining the potential of mechanisms linked to the regional dimension, we decompose the aggregate labor share decline into within- and between-region components. This is analogous to our earlier within/between industry decomposition, where the subscript j in Equation (17) refers to one of the nine Census divisions.

Panel B. in Table 4 displays the results. As with the industry-level exercise, most action occurs *within* regions rather than reflecting between-region reallocation: of the 7.3 ppt decline per decade between 1982 and 2007, 6.6 ppt occur within Census divisions, whereas between-division reallocation accounts for less than a percentage point, even when adding the residual term. An analogous analysis at the state level shows similar results.[20]

5.2.3   Types of firms

Legal form of organization   We next turn our attention to the legal form of organization (LFO) of the firm. This dimension may be relevant if firms with a specific legal form of organization lower their labor share more than their peers. For example, there exists an extensive body of work on the impact of the 1986 tax reform and how it gave rise to a new legal form of organization of businesses: the S-corporation. Smith, Yagan, Zidar, and Zwick (2017) finds that pass-through businesses such as S-corporations, which have accounted for an increasing share of economic activity, are important drivers of income inequality because a few highly profitable mid-sized firms passes through a large portion of business income to the firm owner.

More generally, we can broadly identify the potential role played by the various types of legal form of organization (corporations, proprietorships, partnerships, co-operatives, etc.) by decomposing the aggregate labor share decline into within- and between-type margins, analogously to our previous exercise at the industry and regional levels (j in Equation (17) now refers to a type of legal form of organization). Panel C. in Table 4 reports the results.

Based on our sample, we find that S-corporations in the manufacturing sector display on average a higher labor share than their peers. However, our shift-share decomposition exercise does not identify a substantial role for reallocation across legal forms of organization. While this finding may be different in other sectors where entrepreneurship is more relevant, it does make us skeptical that this dimension is an important factor behind the decline in the aggregate labor share.

Public versus private firms   Lastly, we contrast publicly-traded and privately-held firms. This is of interest for two reasons: First, firms that are publicly traded are less likely to be financially constrained than private firms, potentially allowing them to more easily build new capital and bring down their labor share. Second, since the 1980s, private equity firms and hedge funds have

---

[20]Estimating if establishments are more likely to become hyper-productive once the state enacts right-to-work legislation, we find a statistically significant but economically small effect.

from time to time been very active in buying out public firms. If the new owners are successful in restructuring the activites of the newly-acquired entities in order to increase their productivity, this may drive the overall labor share lower. In the context of the general empirical strategy employed in this section, both theories should show up in the reallocation term of a shift-share decomposition applied to private and public firms. Panel D. in Table 4 shows that this is not the case.

6   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 establishment-level data for the U.S. manufacturing sector between 1967 and 2012. 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 establishment *rose* over the same time period. This apparent disconnect is due a drastic reallocation of production from high-labor share establishments towards their low-labor share peers, which we label hyper-productive, HP, establishments.

We then highlight a number of striking micro-level empirical facts that are difficult, as a whole, to reconcile with the leading theories that have been proposed in the literature. Among others, we show that HP establishments (1) have highly transient labor shares; (2) are more and more likely to have been large before earning their hyper-productive status; (3) do not pay higher wages than their peers; and (4) achieve higher nominal productivity in a large part through higher prices. Our counterfactual exercises indicate that selection along size is the most likely driver behind the fall in the labor share.

These findings, taken as a whole, provide a guide for researchers intent to understand and model the forces that underlie the decline in the aggregate labor share in particular, and establishment or firm level dynamics in general.

References

Daron Acemoglu and Pascual Restrepo. The race between machine and man: Implications of technology for growth, factor shares and employment. *American Economic Review*, 108(6):1488–1542, June 2018. (Cited on page 4.)

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. (Cited on page 4.)

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. (Cited on page 4.)

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. (Cited on pages 4 and 11.)

David Rezza Baqaee and Emmanuel Farhi. Productivity and misallocation in general equilibrium. *Working Paper*, 2017. (Cited on page 5.)

Simcha Barkai. Declining labor and capital shares. *Working Paper*, 2017. (Cited on page 5.)

Daniel Berkowitz, Hong May, and Shuichiro Nishioka. Does capital-labor substitution or do institutions explain declining labor shares? *Working Paper*, 2017. (Cited on page 11.)

Andrew B. Bernard, Stephen J. Redding, and Peter K. Schott. Multiple-product firms and product switching. *American Economic Review*, 100(1):70–97, March 2010. (Cited on page 27.)

Andrew B. Bernard, Stephen J. Redding, and Peter K. Schott. Multiproduct firms and trade liberalization. *Quarterly Journal of Economics*, 126(3):1271–1318, August 2011. (Cited on page 27.)

Christoph E. Boehm, Aaron Flaaen, and Nitya Pandalai-Nayar. Multinationals, offshoring, and the decline of U.S. manufacturing. *Working Paper*, 2015. (Cited on page 4.)

Benjamin Bridgman. Is labor's loss capital's gain? gross versus net labor shares. *Working Paper*, 2014. (Cited on page 5.)

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. (Cited on page 5.)

Russell W. Cooper, John C. Haltiwanger, and Jonathan L. Willis. Declining dynamism at the establishment level. *SED Meeting Paper*, 2017. (Cited on page 25.)

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. (Cited on page 25.)

Ryan A. Decker, John Haltiwanger, Ron S. Jarmin, and Javier Miranda. Changing business dynamism and productivity: Shocks vs. responsiveness. *NBER Working Paper No. 24236*, 2017b. (Cited on pages 25 and 35.)

Maya Eden and Paul Gaggl. On the welfare implications of automation. *Review of Economic Dynamics*, 29:15–43, July 2018. (Cited on pages 4 and 6.)

Chris Edmond, Virgiliu Midrigan, and Daniel Y. Xu. How costly are markups? *NBER Working Paper No. 24800*, 2018. (Cited on pages 5 and 28.)

Jan Eeckhout and Jan De Loecker. The rise of market power and the macroeconomic implications. *NBER Working Paper No. 23687*, 2017. (Cited on page 5.)

Andrea Eisfeldt, Antonio Falato, and Mindy X. Zhang. Human capitalists. *Working Paper*, 2018. (Cited on page 46.)

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. (Cited on pages 4, 5, and 6.)

Lucia Foster, John C. Haltiwanger, and Chad Syverson. Reallocation, firm turnover, and efficiency: Selection on productivity or profitability? *American Economic Review*, 98(1):394–425, March 2008. (Cited on pages 20, 27, and 47.)

Jason Furman and Peter Orszag. A firm-level perspective on the role of rents in the rise in inequality. *Working Paper*, 2015. (Cited on page 4.)

Andrew Glover and Jacob Short. Demographic origins of the decline in labor's share. *Working Paper*, 2018. (Cited on page 4.)

Emilien Gouin-Bonenfant. Productivity dispersion, between-firm competition and the labor share. *Working Paper*, 2018. (Cited on page 11.)

Gene M. Grossman, Elhanan Helpman, Ezra Oberfield, and Thomas Sampson. The productivity slowdown and the declining labor share: A neoclassical exploration. *Working Paper*, 2017. (Cited on page 4.)

Gustavo Grullon, Yelena Larkin, and Roni Michaely. Are US industries becoming more concentrated? *Working Paper*, 2016. (Cited on page 4.)

Germán Gutiérrez and Thomas Philippon. Investmentless growth: An empirical investigation. *Brookings Papers on Economic Activity*, 2017(1):89–169, 2017. (Cited on page 26.)

Barney Hartman-Glaser, Hanno N. Lustig, and Mindy X. Zhang. Capital share dynamics when firms insure managers. *Journal of Finance*, forthcoming. (Cited on page 4.)

Hugo A. Hopenhayn. Entry, exit, and firm dynamics in the long run equilibrium. *Econometrica*, 60(5):1127–1150, September 1992. (Cited on page 13.)

Chang-Tai Hsieh and Peter J. Klenow. Misallocation and manufacturing TFP in China and India. *Quarterly Journal of Economics*, 124(4), November 2009. (Cited on page 20.)

Cosmin Ilut, Matthias Kehrig, and Martin Schneider. Slow to hire, quick to fire: Employment dynamics with asymmetric responses to news. *NBER Working Paper No. 20473*, September 2014. (Cited on page 25.)

Boyan Jovanovic. Selection and the evolution of industry. *Econometrica*, 50(3):649–670, May 1982. (Cited on page 13.)

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. (Cited on page 2.)

Loukas Karabarbounis and Brent Neiman. The global decline of the labor share. *Quarterly Journal of Economics*, 129(1):61–103, February 2014a. (Cited on pages 4 and 8.)

Loukas Karabarbounis and Brent Neiman. Capital depreciation and labor shares around the world: Measurement and implications. *Working Paper*, 2014b. (Cited on page 5.)

Barış Kaymak and Immo Schott. Corporate tax cuts and the decline of the labor share. *Working Paper*, 2018. (Cited on page 4.)

Matthias Kehrig. The cyclicality of productivity dispersion. *US Census Bureau Center for Economic Studies Paper No. CES-WP-11-15*, May 2011. (Cited on pages 4 and 26.)

Dongya Koh, Raül Santaeulàlia-Llopis, and Yu Zheng. Labor share decline and intellectual property products capital. *Working Paper*, 2016. (Cited on pages 4 and 6.)

Robert Z. Lawrence. Recent declines in labor's share in US income: A preliminary neoclassical account. *NBER Working Paper No. 21296*, 2015. (Cited on page 4.)

Brent Neiman and Joseph Vavra. The rise in household spending concentration. *Working Paper*, 2018. (Cited on page 5.)

Ezra Oberfield and Devesh Raval. Micro data and macro technology. *NBER Working Paper No. 20452*, 2014. (Cited on page 4.)

Matthew Rognlie. Deciphering the fall and rise in the net capital share: Accumulation or scarcity? *Brookings Papers on Economic Activity*, Spring 2015. (Cited on page 6.)

Matthew Smith, Danny Yagan, Owen Zidar, and Eric Zwick. Capitalists in the twenty-first century. *Working Paper*, 2017. (Cited on page 41.)

T. Kirk White. Recovering the item-level edit and imputation flags in the 1977-1997 Census of Manufactures. *Census Discussion Paper CES 14-37*, September 2014. (Cited on page 48.)

T. Kirk White, Jerome P. Reiter, and Amil Petrin. Imputation in U.S. manufacturing data and its implications for productivity dispersion. *Review of Economics and Statistics*, 100(3):502–509, July 2018. (Cited on pages 18, 27, and 47.)

Appendix

A   Data and measurement

A.1   Constructing the Full Census Sample

The data used in this project are compiled by the U.S. Census Bureau and comprise the Census of Manufactures (CMF) and – for robustness checks – the Annual Survey of Manufactures (ASM). They are both mail-back surveys and cover the U.S. manufacturing sector (NAICS 31-33) at the establishment level, where an establishment is defined as a distinct unit of a manufacturing firm where the predominant activity is production. Data are collected in 1963 and subsequently in years ending in 2 and 7 since 1967. Some key variables on labor compensation are missing in the 1963 Census, so we drop that year.

In principle, the Census covers all existing 300-350k establishments in the manufacturing sector. We only consider those establishments in the "tabbed sample," a distinction Census started to record in 2002. Non-tabbed establishments are considered by Census to be not really active or are only based on administrative records and thus excluded from publicly available tabulations (hence the name "tabbed"). We follow Census in their assessment of these establishments as not really contributing to economic activity and drop them.

The data carry a wide array of variables only some of which are of interest for this project. These are data on sales, inventories, intermediate and energy inputs, employment and hours, salaries, wages and ancillary labor costs, capital stocks and investment. The following sections describe how observed variables are used to construct measures needed for our analysis. In principle, the labor share is the ratio of total labor costs (described in Section A.2) and value added (described in Section A.3).

A.2   Measuring labor compensation

Labor costs in the Census data consist of three parts: salaries and wages (item SW) which comprise both wages of production workers as well as the salaries of non-production workers. Production workers comprise employees up to and including the line-supervisor level engaged in the core manufacturing activities such as fabricating, processing, assembling, inspecting, receiving, packing, warehousing, maintenance, repair, janitorial and guard services and record keeping. Non-production workers, in contrast, are employees above line-supervisor level which comprises executive, purchasing, professional and technical sales, logistics, advertising, credit, clerical and routine office functions. The third portion are ancillary labor cost, which can broadly be interpreted as benefits. Benefits contain involuntary labor costs (item ILC) such as such as mandatory state pension fund contributions, unemployment insurance, or social security contributions netted out from wages. Voluntary labor costs (item VLC) comprise health, additional voluntary retirement contributions and other benefits paid to employees.

What is missing from labor compensation is compensation in assets such as stock options. While that type of compensation is taxed as labor income when the option is exercised, it is not recorded as labor compensation when the stock option is given. Though this is likely to bias our labor cost and thus our labor share measure downward, we think that bias is small given that only executives are given stock options.[21] Another portion of labor income that is missing is proprietary income. If a lot of the labor share decline was due to more and more labor compensation for entrepreneurs

---

[21]Ongoing research in finance is concerned with the rising share of deferred compensation in total labor compensation, see Eisfeldt, Falato, and Zhang (2018).

funneled as income, we would likely see a strong difference in the labor share by legal form of organization. In particular, we would expect a stronger decline of the labor share for private firms, or "S corporations." This is, however, not the case in manufacturing. We conclude that neither stock options nor proprietary income are a likely cause of the aggregate labor share decline.

A.3 Measuring value added

Value added in the Census data is measured as sales (item `TVS`) less inventory investment for final (difference between `FIE` and `FIB`) and work-in-progress goods (difference between `WIE` and `WIB`), resales (item `CR`)[22], material inputs (sum of items `CP`, `CW` and `MIB` less `MIE`) and energy expenditures (sum of items `CF` and `EE`). Unlike in the BLS data, purchased services is only imperfectly captured in contract work (item `CW`). To account for that, we use the industry-year-specific share of purchased services in sales to deflate sales and then subtract all measured intermediate inputs (CP+CW+(MIB-MIE)+CF+EE).

A.4 Constructing the Matched Price Sample

We are grateful to Kirk White from the U.S. Census Bureau for aiding with the Product Trailer, especially with the edit-in flags.

We combine the product trailers to the Census of Manufactures into a panel of close to nine million product-establishment-year observations. Of these, we keep only observations, in which the variables product value shipped (item `PV`) and product quantity shipped (item `PQS`) are populated and where the latter variable has a meaningful interpretation, say short tons of aluminum sheets or cubic feet of liquefied gas rather than number of vehicles. Census defines a product based on a 10-digit code whose first six digit refer to the 6-digit NAICS industry code. With each of these industries, Census provides a detailed definition of products about which firms have to report product-level sales and – when applicable – the physical quantity produced and shipped.

Only about 130 thousand year-establishment-product observations have that information; similar to the procedure in Foster, Haltiwanger, and Syverson (2008), even though these authors limit attention to 6-digit NAICS industries with homogeneous products, we consider a broader set of multi-product establishments, as long as these products have a well-defined notion of quantity (metric tons of chemicals, ...)

In addition to that, we limit attention to observations that are not imputed in a way that would change the empirical variance of the `PV` or the `PQS` distributions. Census uses an array of criteria to delete originally reported data when they fail certain reasonability tests. These values are then replaced by imputed data where an algorithm chooses from about a dozen different imputation methods the one which mostly likely replicates the correct aggregates. White, Reiter, and Petrin (2018) have developed an improved method that changes imputations to not only correctly replicate aggregates, but also preserves the cross-sectional distribution. We have not obtained their toolbox yet, but plan to do so in the future. This means that for now, we have to rely on observations that are not imputed in a way that would change the cross-sectional distribution. These are labeled by the following edit-in flags that consist of three letters:

- `R__`: Any observation starting with `R` denotes reported values. Of these, we keep those that were not replaced with an imputed value, in particular:

    - `RC`: analyst correction of reported value,

---

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

- `RG`: goldplated observation (due to analyst information "known" to be of such high quality that any imputation would worsen data quality),

- `RN`: reported value just corrected for obvious rounding errors;

- `RO`: override imputation with establishment-specific information (say, information obtained in a phone call);

- `RU`: preserve reported value due to inability to perform imputation;

- `RZ`: reported zero which is acceptable.

- `_C_`: any observation with `C` in the middle – whether originally reported (observations that start with `R`) or not reported and then filled in by information through other means such as follow-up phone calls (observations that start with a blank value) – refers to values that have been corrected by an analyst using establishment-specific information.

- Observations that start with a `C` should not occur according to the Census system of edit-in flags. We assume that the roughly twenty thousand observations in 1992 and 1997 are erroneously coded and mean to start with a blank and should be `_C`.

One limitation of that approach is that we are constrained to data since 1992 as observations in the product trailer do not carry edit-in flags prior to that year. White (2014) has recovered these flags from the raw datafile that are not accessible to RDC researchers at this point, but we hope to obtain them in the future, so we can extend our analysis back to 1977. At this point, we are left with about 130 thousand usable and non-imputed product-year-establishment prices which aggregate up to about 41 thousand establishments, so the typical establishments produces and sells on average a bit more than three products. Prices at the 10-digit NAICS product level are finally constructed by dividing `PV` by `PQS`.

A.5   Comparison Full Census Sample vs. Matched Price Samples

We study the differences in sales per worker between the Full Census Sample and the Matched Price Sample in which we observe product prices and quantities separately. The objective is to show that in the Matched Price Sample, the same cross-sectional patterns of sales per worker vis-à-vis the labor share and the dynamic differences of sales per worker growth between HP and Non-HP establishments exist.

In order to produce Figure A.1, we run a non-parametric regression analogous to Equation (9) of relative sales per worker on the labor share in both the Full Census Sample and the more homogeneous Matched Price Sample. Even though the relative differences of sales per worker might not be as pronounced in the latter, the relationship between relative sales per worker and the labor share look very similar across the two samples. Only at very low labor shares are sales per worker in the Matched Price Sample significantly lower than those in the full panel, but the differences with other establishments remain stark. For example, establishments with a labor share of 10 ppt still generate generate 1.7 times ($\exp(0.53) \approx 1.7$) more sales with the same workforce than the average establishment. In the Full Census Sample this number is 2.3.

In the right panel of Figure A.1 we display the relative sales-per-worker dynamics of HP establishments versus Non-HP establishments. The approach is analogous to (3) and (4): we regress the growth rate of sales per worker, $\Delta(\widetilde{pq/l})$, on a dummy variable that equals one if establishment i is an HP establishment. This regression is done in both the Full Census and the Matched Price Sample, with the intention of studying how much the sales-per-worker dynamics differ the in the two samples. In the Full Census Sample, sales per worker of HP establishments jump relative to

Figure A.1: Relative sales per worker in the Full Census Sample vs. the Matched Price Sample

Note: The left panel in the figure depicts the cross-sectional differences in relative sales per worker $\widetilde{pq/l}$ against the labor share in the Full Census Sample (dark grey line) and the Matched Price Sample (light grey line). Dashed lines denote 95% error bands.
The right panel displays the cumulative growth of relative sales/worker $\Delta(\widetilde{pq/l})$ of HP establishments in both samples. Whiskers denote 95% error bands.

the Non-HP establishments by 21% during the five years preceding the year in which they become HP. In the subsequent five years, more than two thirds of that relative sales growth is erased and the 10-year differential growth rate is only 6.7% more for HP vs Non-HP establishments. Over the entire time span, the estimates for the Full Census Panel show a significantly different sales per worker trajectory for HP establishments than for Non-HP establishments.

The evidence in the Matched Price Sample exhibits a similar qualitative pattern. Unsurprisingly, the magnitudes are smaller because the establishments in the Matched Price Sample are much more homogeneous than in the Full Census Sample. In the five years preceding an establishment's HP status, sales per worker grow by 12.5% more for HP establishments and revert to about 5% in the subsequent five years. Due to the smaller sample, these estimates are noisier for the Matched Price sample.

B   Robustness

In this appendix, we carry out some robustness check about our empirical findings.

B.1   Transitory versus permanent HP establishments

In Section 3.2 we showed that HP establishments are largely a temporary phenomenon and that their labor shares display a V-shaped pattern in the years surrounding the time they are in the lowest quintile of labor shares in a given industry. Obviously, some of the HP establishments do have a permanently low labor share and are among the HP establishments for several Census years in a row, while others display an even more volatile labor share. We want to understand the role of "permanent" versus "transitory" HP establishments. Since the former tend to be larger and thus more relevant for aggregates, we want to ensure that the "temporary HP establishments," those characterized by the V-shaped pattern of Figure 8, play a significant role for the aggregate labor share decline.

To that end, we partition the set of HP establishments in period t into those that are an HP establishment from $t-4$ to $t+5$, denoted "permanent HP," and the rest, denoted "temporary HP." When we drop both temporary and permanent HP establishments from the sample, the aggregate labor share has a much higher level and stagnates. This shows that HP establishments are essential to understanding the aggregate labor share decline; see the light grey line in Figure B.2. When we instead only drop the permanent HP establishments, however, the counterfactual labor share dynamics do not look markedly different: while the *level* is somewhat higher by definition (these are, after all, low-labor-share establishments), the overall *decline* is similar in magnitude to that of the actual labor share. This confirms that temporary HP establishments play an important role for the aggregate labor share level and its decline.

Figure B.2: The role of temporary and permanent HP establishments

Aggregate labor share with and without HP establishments

B.2 Firms or establishments?

In this section, we study the labor share at the level of the firm. Two considerations motivate this analysis. First, we showed that price dynamics are responsible for a large share of sales-per-worker and labor-share dynamics at the establishment level. If these prices are transfer prices across establishments within the same firm rather than market sales prices, the labor share of firms will likely be much more smooth regardless of their labor share level. Second, if the price and productivity drivers of the labor share derive from firm factors such as brand power or superior management practices, then HP establishments likely sort into the same firms. Labor shares of firms that operate mostly HP establishments should then exhibit the same V-shaped pattern that we observe for the HP establishments in Figure 8. If, on the other hand, HP establishments are evenly distributed across firms, we would expect firm-level labor shares to be much more stable and as establishment-level labor share dynamics get diversified away by the firm.

To that end, we aggregate labor cost and value added across all establishments within the same firm (defined by `FIRMID`) to compute firm-level labor shares. Analogously to HP establishments, we define "HP firms" as firms whose labor share is in the lowest quintile of their modal industry in a given year. We then repeat the analysis of (3) and (4) for these HP firms and show them in Figure B.3. Clearly, the V-shaped pattern is still present at the firm level even though the magnitude is slightly smaller for the weighted estimate. For the unweighted (not displayed), the V-shapes look equally large. This leads us to two conclusions: First, within-firm transfer prices are not the main driver of the price dynamics documented in Section 3.6. Second, HP establishments tend to assort

50

into the same firms.

Figure B.3: HP establishments versus HP firms

Change of labor share of HP establ./firms (in ppt): CMF

Cumulative change from ...

$t-5$ to t $t-5$ to $t+5$

HP establishments
HP firms