The Micro-Level Anatomy of the Labor Share Decline*

Matthias Kehrig† and Nicolas Vincent‡

December 2019

Abstract

The manufacturing 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 salient facts: (1) since the 1980s, there has been a dramatic reallocation of value added towards the lower end of the labor share distribution; (2) this aggregate reallocation is not due to entry/exit, "superstars" growing faster or large establishments lowering their labor shares, but instead by units whose labor share falls as the same time as they grow in size; (3) low-labor-share (LL) establishments have only temporarily lower labor shares that rebound after five to eight years to the level of their peers; (4) low labor shares are driven by high revenue labor productivity, not low wages; (5) LL establishments enjoy a product price premium relative to their peers that causes their high (revenue) productivity, pointing to a significant role for demand-side forces; (6) there is evidence that advertising activity plays a role in shaping labor share dynamics; (7) the transient pattern of labor shares for LL establishments has become more pronounced over time; and (8) the dynamics of value added and employment have become more and more disconnected over time.

Keywords: Labor Share, Productivity, Firm Size Distribution, Relative Prices.
JEL classification: E2, L1, L2, L6, O4.

*We would like to thank Gadi Barlevy, Nick Bloom, Jeff Campbell, Julieta Caunedo, John Cochrane, Allan Collard-Wexler, Steve Davis, Hugo Hopenhayn, Chang-Tai Hsieh, Chad Jones, Pete Klenow, Brent Neiman, Luigi Pistaferri, Peter Schott, Chad Syverson, T. Kirk White, Daniel Xu and our discussants Zsofia Barany, Chris Gust, Ezra Oberfield and Kirill Shakhnov as well as seminar participants at Chicago-Booth, BFI, Rochester, Houston, Columbia, BI, CERGE-EI, Clemson, Stanford, Penn State, Mannheim, ITAM, UVA, UAB, Royal Holloway, UCL, UNC, LMU Munich, UIUC, Cornell, Auburn, the Federal Reserve Banks of Chicago, St. Louis and Philadelphia, the 2015 meeting of German Economists Abroad, the 2016 Cambridge-INET conference "Firms in Macroeconomics," the 2016 Barcelona Summer Forum, the 2017 Annual Meeting of the SED, the 2017 UniCredit Alumni Workshop, the Société Canadienne de Science Économique, the 2018 Canadian Economics Association Meetings, the 2018 ERID Conference on "Firm, Industry, and Trade Dynamics," the 2018 Northwestern Macro Alumni Meeting, the ESSIM 2018, the 2019 ASSA meetings, the 2019 Goldwater Conference at ASU, for helpful comments about earlier version of this project. All errors are our own. We thank Xian Jiang and Vytautas Valaitis for excellent research assistance. Financial support from the Fondation HEC Montréal (Vincent) and the National Science Foundation under NSF grant No. SES-1758426 (Kehrig) are gratefully acknowledged. Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed.

†Duke University, Department of Economics, NBER and CEPR. Email: matthias.kehrig@gmail.com.
‡HEC Montréal, Department of Applied Economics. Email: nicolas.vincent@hec.ca.

Contents

List of Tables

List of Figures

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 and which shocks are consistent with them. 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 salient facts and interpret their implications in the context of a simple conceptual framework. 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 sectoral 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 other words, the covariance between establishment-level labor shares and value added weights turned sharply negative over time. In contrast, a similar reallocation of labor was almost non-existent over the same period.

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

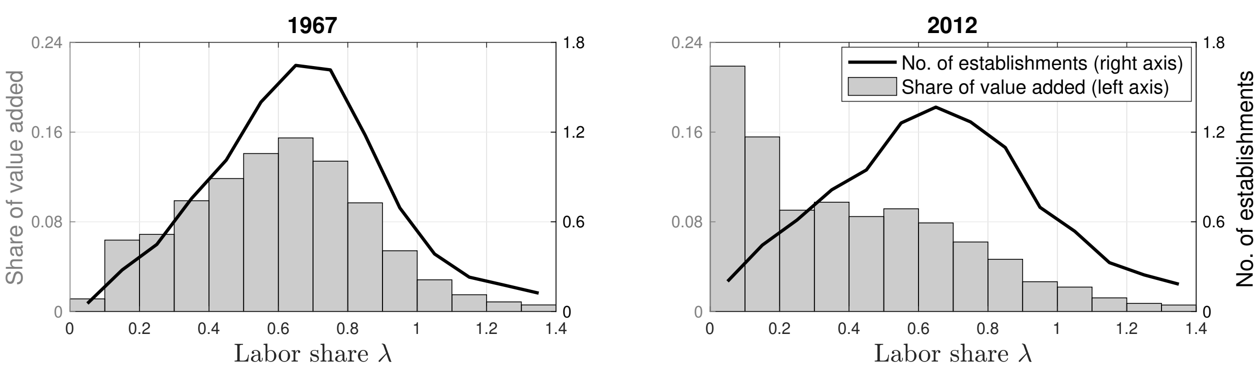

Note: The solid black line (right axis) reflects the raw cross-establishment distribution of labor shares, while the distribution of value added is represented by the grey bars (left axis). See notes to Figure 3 for more details.

Next, we investigate the reasons behind the decline in the covariance between labor share and value added weight at the establishment level. We show that this was not driven by entry/exit nor by large establishments lowering their labor shares over time nor by the rise in economic importance of low-labor-share "superstars". Instead, we establish that the strong reallocation was driven by units whose labor shares fell *at the same time as they grew in size.*

Third, we find that labor share fluctuations at the establishment level are surprisingly transient: only about 8% to 18% of the deviation of an establishment's labor share remains after five years. We then focus on low-labor-share (LL) establishments, defined as those in the lowest quintile of the labor share distribution for a given year and industry. As was showed earlier, these units experienced a drastic increase in their economic importance. Yet, LL status is temporary : the probability that a typical LL establishment today loses that status five years later is about 60%, while that number would be close to 0% if LL establishments had permanently low labor shares. Even more strikingly, we document that the labor share of the typical LL establishment follows a pronounced V-shaped pattern over time: the drop in labor share they experience in the five years preceding LL status is almost equal to the rebound in the following five years.

Fourth, we decompose the labor share into its components to understand the main drivers behind the micro-level dynamics. We find that the cross-sectional labor share differences are almost entirely driven by differences in value added. Similarly, the V-shaped pattern of the labor share of LL establishments is mainly driven by fluctuations in value added, with very little accounted for by employment or wages.

Fifth, we use a subsample of the Census database which provides information about the value of sales and physical 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 LL establishments.

What do we learn from these five findings? When interpreted through the lens of our conceptual framework, they argue for a significant role played by demand factors. While technology shocks could in theory also generate a negative comovement between labor share and value added if paired with endogenous markups, they would counterfactually predict that prices *drop* with labor shares. Similarly, the lack of a relationship between wages and labor shares, both in levels and dynamics, cast doubts on explanations centered around a rise in monopsony power in labor markets.

We complete our empirical analysis by highlighting additional findings that deepen or confirm our main conclusion that demand shocks are driving labor shares and value added shares. Using data on advertising expenditures, we document a strong positive relationship between advertising intensity and labor shares. Given the role of advertising in shaping customer behavior, we see this evidence as supportive of the role bolsters our conclusion that demand factors play a significant role. Furthermore, we find that the effect of advertising intensity on subsequent value added growth strengthened over time. This is consistent with evidence that the V-shaped pattern of the typical

LL-establishment's labor share deepened over time, alongside a stronger disconnect between value added and employment dynamics.

The paper is organized as follows. The next section discusses how our paper fits in the recent labor share literature. In Section 3, we present a simple conceptual framework to guide the interpretation of the empirical findings from Sections 4 and 5. We conclude in Section 6.

2 Relation to the literature

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 (2019) 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, Böckerman and Maliranta (2012) present evidence that exposure to international trade is related to the labor share decline in Finland. Elsby, Hobijn, and Şahin (2013) advocate the role of offshoring as an important driver of the labor share decline in the U.S. 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 (2019) 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 "superstar establishments" are mainly responsible for the manufacturing labor share decline. Grullon, Larkin, and Michaely (2016) use firm-level data from Compustat to document that most U.S. industries became more concentrated over time, with the "winning firm" making large profits and realizing outstanding stock returns as well as more profitable mergers and acquisitions. Barkai (2017) and De Loecker, Eeckhout, and Unger (2018) 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 (forthcoming) 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 (2019) 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.

3 Conceptual framework

The main objective of this paper is to study the micro-level anatomy of the aggregate labor share decline. Many different causes – patterns of reallocation across micro units, different types of shocks – may lead to that outcome. Knowing which causes hold empirical ground will help us understand those structural features of the U.S. economy that matter for the labor share decline. In this section, we lay out a succinct conceptual framework to guide our analysis, built around a simple production function. Its purpose is not to undertake a formal quantitative assessment of different causes, but to identify the qualitative relevance of a variety of shocks and reallocation dynamics that could be behind the aggregate labor share decline. Throughout the empirical analysis , we will often refer back to this conceptual framework to interpret our findings.

3.1 The aggregate labor share

To frame our analysis, consider a specific production unit i (firm, plant, etc.) at time t that employs L_{it} workers at wage rate W_{it} to produce Y_{it} units of a good sold at price P_{it}. The labor share of that unit is then the ratio of its labor cost to the nominal value added: $\lambda_{it} \equiv (W_{it}L_{it})/(P_{it}Y_{it})$. Summing up across units, one can express the aggregate labor share, λ_t, as the weighted sum of the individual labor shares:

$$\lambda_t = \frac{W_t L_t}{P_t Y_t} = \frac{\sum_i W_{it} L_{it}}{\sum_i P_{it} Y_{it}} = \sum_i \omega_{it} \lambda_{it} \tag{1}$$

$$= \overline{\lambda}_{it} + Cov(\lambda_{it}, \omega_{it}) \tag{2}$$

where λ_{it} corresponds to the labor share of production unit i at time t and $\omega_{it} \equiv \frac{P_{it}Y_{it}}{\sum_i P_{it}Y_{it}}$ denotes the value-added weight of unit i. The second line in the above expression derives from Olley and Pakes (1996) and is useful to illustrate how the aggregate labor share depends on the common unweighted average, $\overline{\lambda}_{it}$, as well as the joint distribution of labor shares and value added, $Cov(\lambda_{it}, \omega_{it})$. We now turn our attention to two broad types of distributional changes that are compatible with an aggregate labor share decline. Then, in the next section, we present a set of candidate economic shocks at the micro-level that can rationalize these changes.

Common effect/trend First, a fall in the aggregate labor share can be the result of a decline in the unweighted average of the distribution of labor shares. That is, any change that affects all or most units symmetrically will alter $\overline{\lambda}_{it}$.

Composition effects Second, equation (2) indicates that the aggregate labor share can decline if the joint distribution of labor shares and value added shares evolves in a way that reduces the covariance between these two objects. This can happen for three reasons:

1. Changes in unit-level market share, $\Delta\omega_{it}$, may be negatively correlated with the initial level of labor shares λ_{it-1}. For example, "superstar units" with constant but lower-than-average labor shares may be growing faster over time. As a result the covariance term in equation (2), $Cov(\lambda_{it}, \omega_{it})$, would decline because $Cov(\lambda_{it-1}, \Delta\omega_{it}) < 0$.

2. Conversely, labor share changes, $\Delta\lambda_{it}$, may be correlated with initial size ω_{it-1}. For example, large units could be more successful in lowering their wages while keeping employment

and output constant, in turn depressing their individual labor shares. The covariance term, $Cov(\lambda_{it}, \omega_{it})$, would fall since $Cov(\Delta\lambda_{it}, \omega_{it-1}) < 0$.

3. Lastly, labor share changes, $\Delta\lambda_{it}$, and relative growth, $\Delta\omega_{it}$, may be negatively correlated. For example, some units may experience shocks or take actions that lead them to simultaneously gain market share and lower their labor share, reducing the $Cov(\lambda_{it}, \omega_{it})$ term.

The discussion in this section makes it clear that the micro-level dynamics of labor shares and market shares can impact the aggregate labor share through many channels. Next, we identify a number of micro-level shocks that may shape these dynamics through their impact on the components of labor and market shares: wages, employment, prices and real output.

3.2 Micro-level effects of demand, supply and monopsony shocks

From de-unionization to automation to rising market power, different forces may impact labor shares at the micro level through distinct components such as wages or markups. In the empirical section, we will study those components, with the aim of identifying explanations that are less likely to be relevant and theories that merit attention in further research. To frame our analysis, consider that the production unit i takes factor prices as given, hires labor L_{it} and rents capital K_{it} to produce output Y_{it} using a Cobb-Douglas production function: $Y_{it} = A_{it}K_{it}^{\alpha_i}L_{it}^{1-\alpha_i}$ where $\alpha_i \in (0,1)$.[1] Unit i's labor share can be written as:

$$\lambda_{it} = \frac{W_{it}L_{it}}{P_{it}Y_{it}} = \frac{W_{it}}{ARPL_{it}} = \frac{W_{it}}{P_{it}APL_{it}} \tag{3}$$

where W_{it} denotes the market wage, while $ARPL_{it} = \frac{P_{it}Y_{it}}{L_{it}}$ and $APL_{it} = \frac{Y_{it}}{L_{it}}$ are the average revenue and physical products of labor, respectively. Next, we analyze, through the lens of this simple framework, three broad classes of theories that have been proposed in the literature to explain the decline in the labor share.

Demand shocks and markups Let us decompose further both λ_{it} and ω_{it}:

$$\lambda_{it} = \frac{W_{it}}{P_{it}APL_{it}} = \frac{W_{it}}{\mu_{it}MC_{it}APL_{it}} = \frac{1-\alpha_i}{\mu_{it}} \tag{4}$$

$$\omega_{it} = \frac{P_{it}Y_{it}}{\sum_i P_{it}Y_{it}} = \frac{\mu_{it}MC_{it}Y_{it}}{\sum_i P_{it}Y_{it}} \tag{5}$$

where μ_{it} corresponds to the markup of the price, P_{it}, over the marginal cost, MC_{it}. The last expression for the labor share follows from the Cobb-Douglas production function, where $1-\alpha_i$ corresponds to the labor elasticity of output. Consider that, for some reason, customers value unit

[1] Though we focus on constant returns to scale in this exposition, the subsequent arguments hold for production functions with decreasing returns to scale as well.

i's products or brand image more so than that of the competition. With an isoelastic demand schedule, the only impact of this preference shock would be to raise the unit's market share, ω_{it}. The aggregate labor share could in turn be impacted through a composition effect: for example, the concentration of preference shocks on low-labor-share units would imply that $Cov(\lambda_{it-1}, \Delta\omega_{it})$ is negative.

Alternatively, unit-level labor shares may be affected if markups are instead endogenous. For example, a demand shock may bring unit i into a less elastic part of its demand curve, leading it to increase its markup. From the two equations, we can clearly see how an idiosyncratic demand shock that raises unit i's markup μ_{it} leads to a fall in its labor share λ_{it}, and a rise in its market share ω_{it}. The sources and consequences of rising markups have been extensively studied recently, see for example Grullon, Larkin, and Michaely (2016); Barkai (2017); Baqaee and Farhi (forthcoming); Gutiérrez and Philippon (2017); Edmond, Midrigan, and Xu (2018); De Loecker, Eeckhout, and Unger (2018); Neiman and Vavra (2019).

TFP shocks Technology is another channel that has been suggested by the literature as a potential driver of the downward labor share trend. With a Cobb-Douglas production function, a positive technology shock lowers the unit's marginal cost, MC_{it}, and increases its average labor productivity, APL_{it}, in a way that these changes exactly cancel each other; under our assumptions the only factors specific to unit i's labor share are its production elasticity α_i and its markup μ_{it}. Therefore, higher TFP on its own does not have a direct impact on the unit's labor share λ_{it}, but it will increase its market share ω_{it}.[2] Standard TFP shocks could lower the aggregate labor share if they are correlated with labor share levels, i.e. $Cov(\Delta\omega_{it}, \lambda_{it-1}) < 0$ as described in the "composition" paragraph above. Examples of these type of shocks include Kaymak and Schott (2018), Alvarez-Cuadrado, Long, and Poschke (2015) or Grossman, Helpman, Oberfield, and Sampson (2017).

TFP shocks may have a different impact if producers do not pass through all the cost savings of a technology shock through lower prices. Instead, they may choose to raise markups μ_{it} because, for example, producing on a larger scale brings them to a less elastic portion of their demand schedule, as in Kimball (1995); Melitz and Ottaviano (2008). This would be in line with the explanation of Autor, Dorn, Katz, Patterson, and Reenen (2017a,b). Under this scenario, equations (4) and (5) imply that an idiosyncratic TFP shock will move unit i's labor share and market share in opposite directions. Examples of these shocks are featured in Grossman, Helpman, Oberfield, and Sampson (2017); Leblebicioğlu and Weinberger (2018); Karabarbounis and Neiman (2014a). Notice that these dynamics are similar to those under the scenario of a demand shock with non-isoelastic demand schedules, except for one important difference: prices will decline after supply-side TFP shocks, while they increase under demand shocks.

[2] This assumes a price elasticity of demand larger than unity as standard in the literature. These points generalize to Cobb-Douglas production functions with non-constant returns to scale and CES production functions with constant returns to scale and Hicks-neutral technology.

Monopsony power Lastly, let us turn to the role of monopsony power in labor markets. If labor market concentration allows businesses to extract rents from workers, we need to relax our assumption that units take factor prices as given. Instead, we follow Berger, Herkenhoff, and Mongey (2019) and rewrite the wage of production unit i, W_{it}, as its marginal revenue product of labor, $MRPL_{it}$, times a generic wage mark-down, denoted $\nu_{it} \leq 1$. The more market power a firm has, the more it can depress the wage beneath the marginal revenue product, which is captured by a lower value of ν_{it}.

$$\lambda_{it} = \frac{W_{it}L_{it}}{P_{it}Y_{it}} = \frac{W_{it}}{ARPL_{it}} = \frac{\nu_{it}MRPL_{it}}{ARPL_{it}} = \nu_{it}(1-\alpha_i) \tag{6}$$

$$\omega_{it} = \frac{P_{it}Y_{it}}{\sum_i P_{it}Y_{it}} \tag{7}$$

Note that by itself, a lower ν_{it} decreases the unit's labor share, but does not increase its value added weight unless it finds it profitable to expand scale and/or adjust its price relative to its peers. A stronger wage mark-down may result from increasing labor market concentration (Azar, Marinescu, and Steinbaum (2017); Hershbein, Macaluso, and Yeh (2018); Berger, Herkenhoff, and Mongey (2019); Jarosch, Nimczik, and Sorkin (2019)); labor market deregulation such as deunionization (Fichtenbaum (2011)); "right-to-work" laws (Blanchard and Giavazzi (2003)); demographic factors (Glover and Short (2018)) or labor market frictions (Gouin-Bonenfant (2018)). While the empirical evidence on the link between business concentration trends and labor shares is ambiguous (Berger, Herkenhoff, and Mongey (2019); Hershbein, Macaluso, and Yeh (2018)), the use of micro-level data will allow us to assess its role for the labor share decline.

This conceptual framework, while simple, provides us with a lens through which we can interpret the micro-level evidence on labor shares, value added, employment, wages and prices. We now turn to documenting a series of empirical findings that inform us on the forces behind the decline in the aggregate labor share.

4 Five findings about the labor share

In this section, we discuss our data sources and describe how we compute the labor share at the manufacturing establishment level. We then produce a number of findings that highlight the micro-level dynamics at play behind the fluctuations of the aggregate labor share.

4.1 Data sources and measurement

Our focus is on the labor share dynamics in the U.S. manufacturing sector. This choice was driven by 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. 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 eroding 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 income 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 the measurement of real estate income: 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 (2019) show that manufacturing is one of the few sectors in which the measured labor share decline is not overturned by the rise in intellectual property products.

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 until 2012.[3] 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. We verify that the labor share dynamics in our Census data coincide with those documented in the Multifactor Productivity Tables published by the Bureau of Labor Statistics (BLS). The labor share λ_t in a given industry and year t is defined as

$$\lambda_t = \frac{W_t L_t}{P_t Y_t} \tag{8}$$

where $W_t L_t$ denotes manufacturing labor costs and $P_t Y_t$ is nominal 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.

We define the following items as labor costs: salaries and wages for permanent (item `SW`) and leased workers (item `CTEMP`), involuntary labor costs (item `ILC`) such as unemployment insurance

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

or social security contributions netted out from wages and voluntary labor costs (item `VLC`) such as health, retirement and other benefits paid to employees.[4] Value added is measured as sales less inventory investment for final and work-in-progress goods, resales, material inputs, contract work, and energy expenditures.[5] 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). For more details on the construction of the sample and the variables of interest, see Appendix A.

Next, we study the anatomy of the decline in the manufacturing labor share by exploiting the establishment-level data described above. We present and analyze five main findings on the micro-level dynamics of the labor share. Our view it that any theory of aggregate labor share dynamics should be compatible with these stylized facts.

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

We start by exploiting the decomposition of the manufacturing labor share λ_t introduced in equation (1) and that we reproduce in the first line below:

$$\begin{aligned} \lambda_t = \frac{\sum_i W_{it} L_{it}}{\sum_i P_{it} Y_{it}} &= \sum_i \omega_{it} \lambda_{it} \\ &= \overline{\lambda}_{it} + Cov(\lambda_{it}, \omega_{it}) \end{aligned} \tag{9}$$

where λ_{it} and ω_{it} correspond respectively to the labor share and value-added weight of establishment i at time t. The second line isolates the role of the covariance: $\overline{\lambda}_{it}$ is the unweighted average and $Cov(\lambda_{it}, \omega_{it})$ the covariance between establishment-level labor and market shares.

From this decomposition, we can readily identify two broad ways that a decline in the aggregate labor share may come about. First, it could follow from a general decline of the unit-level labor shares λ_{it}, which would be reflected in a lower (unweighted average) $\overline{\lambda}_{it}$. This may, for example, come from a rise in markups or monopsony power common to all units. Second, the fall in the manufacturing labor share λ_t could be the result of a decline in the covariance between λ_{it} and ω_{it}. For instance, this would happen if low-labor-share establishments experience an increase in their economic weight over time.

In this section and the next, we aim to disentangle between these various scenarios with the help of micro-level data.

---

[4]Census does not collect information on 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 interested in the rising share of deferred compensation in total labor compensation. This could potentially mitigate the aggregate labor share decline.

[5]Unlike in the BLS data, purchased services are not reported in the Census data. To account for this discrepancy, we reduce establishment-level value added by the industry-year-specific ratio of purchased services to value added computed from the BLS data.

4.2.1 The labor share of the median establishment *increases*

As a first step, Figure 2 plots several quantiles of the raw distribution of establishment-level labor shares λ_{it} in each Census year since 1967, alongside the manufacturing labor share. Figure 2

Figure 2: Sectoral and establishment-level labor shares in U.S. manufacturing

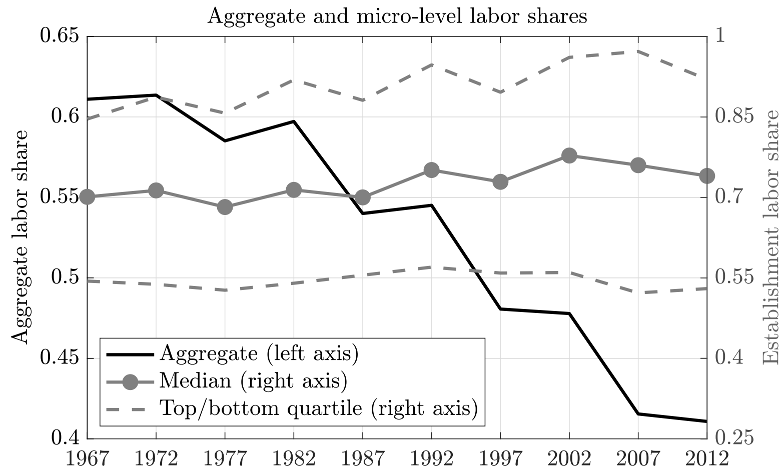

Note: The figure plots the sectoral manufacturing labor share (black line, left axis) against the year-by-year quantiles of the cross-establishment labor share distribution (grey lines, right axis): the solid grey line with balls reflects the median, the dashed grey lines reflect the first and third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

highlights diverging trends in the labor shares at the sectoral and establishment level, particularly since the mid 1980s: while the manufacturing 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. This finding already makes it clear that the manufacturing labor share decline is not mainly driven by a shift of the distribution of labor shares in individual establishments (corresponding to the λ_{it} terms in Equation (1)). Instead, our evidence points to the importance of reallocation (corresponding to the ω_{it} terms in Equation (1)) as the main driver of the manufacturing labor share dynamics.[6] This is what we turn our attention to next.

4.2.2 The dramatic reallocation of value added

The diverging trends between aggregate and micro-level labor shares imply that the ω_{it} terms in Equation (1) must play a central role in driving down the manufacturing labor share, through a reallocation of value added towards the lower tail of the labor share distribution. To quantify this reallocation, we divide the distribution of labor shares λ into 10 ppt-wide bins from 0 to 140 ppt in each year. For each labor-share bin, we then compute its share of total manufacturing value added, employment and number of establishments. To control for industry-specific differences, we

[6] In Appendix B.2, we show that the decline in the manufacturing labor share is present for both production and non-production workers.

compute these shares for each 3-digit NAICS industry and then aggregate them up in each bin using the industry's value-added weight in a given year. The subsequent analysis therefore refers to reallocation of value added *within* a typical industry.[7]

Figure 3: Distribution of establishments, hours worked and value aded along the labor share distribution

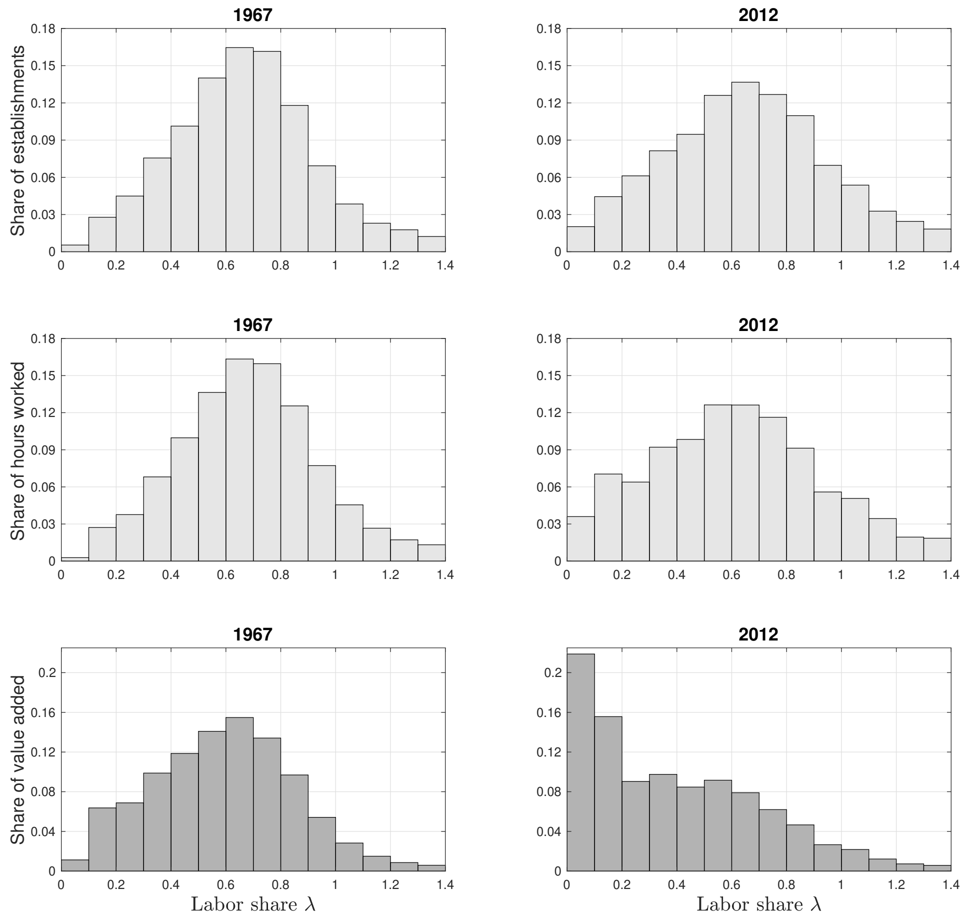

Note: The rows respectively show the distribution of establishments (top), hours worked (middle) and value added (bottom) in 1967 (left) and 2012 (right) against the labor share. The top panels indicate no significant locational shift of establishment-level labor shares from 1967 to 2012. Changes in the distribution of hours worked are also limited (middle row). In both cases there is, however, evidence of a fattening of the tails. The distribution of economic activity (bottom panels), 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.

The top and middle panels of Figure 3 respectively display the distributions of the number of establishments and hours worked against the labor share λ_{it}, in 1967 (left) and 2012 (right). In

[7]Repeating this exercise at other aggregation levels, we find almost no difference between 3-digit and 4-digit NAICS levels, while the reallocation of value added to low labor shares within 6-digit NAICS industries is even stronger.

both cases, the distribution did not change drastically, though there is evidence of some fattening of the tails. The change, however, is much starker for value added, shown in the bottom row. In 1967, most of value added was produced by establishments in the middle of the labor share distribution (between 50 and 80 ppt), with a value-added weighted median labor share of 62 ppt. Over the following decades, however, economic activity shifted drastically towards the low-labor-share spectrum; the bottom-right panel shows that by 2012, half of manufacturing value added is accounted for by establishments with a labor share less than 32 ppt. Similar evidence has been found for other sectors in the U.S. by Autor, Dorn, Katz, Patterson, and Reenen (2017b), for Canada by Gouin-Bonenfant (2018), and for China by Berkowitz, May, and Nishioka (2017). The disconnect between value added and labor reallocation is a key feature of the labor share decline.

Referring back to our discussion surrounding equation (2), Finding 1 makes it clear that common trends, for example due to a generalized increase in markups or monopsony power, are unlikely to be behind the decline in the manufacturing labor share. Instead, it appears to be due to the covariance between labor and market share turning strongly negative: since the 1980s, low-labor-share establishments have also happened to be much larger than their peers. In the next section, we investigate what could be behind this development and argue that the *joint* dynamics of value added and the labor share is central to this phenomenon.

4.3 Finding 2: The importance of the joint labor share and size dynamics

While the evidence in Figure 3 is stark, it does not indicate how the reallocation of value added came about. Consider three very different theories that would generate a fall in the aggregate labor share through a decline in $Cov(\lambda_i, \omega_i)$. First, entry/exit could be behind the aggregate downward trend if the labor share of the typical exiting establishment is increasingly higher than the average, while that of the typical entrant is increasingly lower. Second, differential growth rates can have an aggregate impact if they are correlated to labor share at the establishment level. For example, this would happen if "superstar" units with very low labor shares were to grow at the expense of their peers. Third, large establishments may see their labor shares drop, for example, because of an increase in their markups or monopsony power, while smaller ones experience the opposite trend. All three scenarios would be compatible with the negative covariance between labor and market shares that we documented in Finding 1. In this section, we put them to the test with the help of counterfactuals.

4.3.1 The negligible role of entry and exit

First, we test whether the decline in the aggregate labor share is driven by reallocation of value added through entry and exit: the extensive margin can be quantitatively 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 entry and exit, we compute the manufacturing labor share for a strongly balanced sample[8] of establishments that are permanently active from 1967 to 2012. Figure 4 reveals that the manufacturing labor share in the strongly balanced sample is about 2

Figure 4: Manufacturing labor share in full and strongly balanced sample

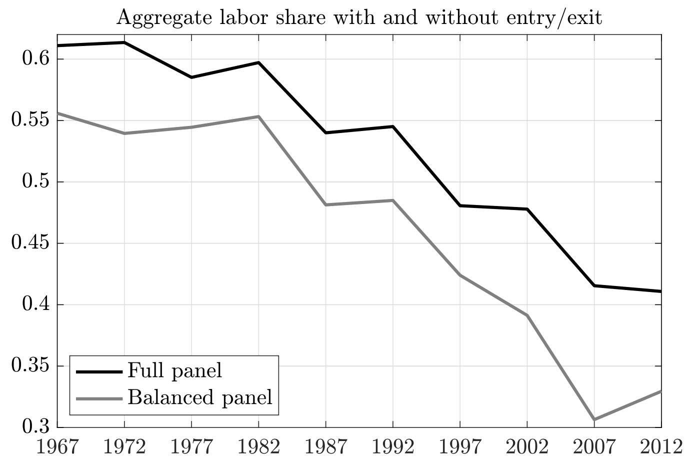

Note: The figure plots the manufacturing labor shares computed on the full panel (solid black line) against that computed on a strongly balanced panel (solid grey line). It shows that entry and exit matter for the *level*, but not for the *decline* of the labor share.

ppt lower than that of the full sample, suggesting that entry and exit indeed depress 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 manufacturing labor share in both samples stagnates until 1982, then falls by 7.3 ppt per decade from 1982 to 2007 in the full sample versus 7.4 ppt in the balanced panel. The takeaway is that the contribution of the extensive margin to the decline of the manufacturing labor share is quantitatively small.[9] Instead, most of the reallocation we documented earlier takes place among incumbent establishments.

4.3.2 Do "big players" or "superstars" drive reallocation?

Next, we turn our attention to the two other scenarios we highlighted earlier. First, we investigate whether establishments that were larger than the average early in the sample (i.e., in the 1970s and 1980s) experienced a relative fall in their labor share afterwards, while smaller establishments instead saw a rise. We refer to this scenario as the "big-player economy"; it could be driven, for example, by the ability of larger employers to keep wages low through monopsony power. Alternatively, establishments with a low initial labor share may have grown faster than their peers

[8] If we focus on the labor share in a sample of long-lived establishments or establishments that were born at least five years ago and will exit at least five year from now, it looks indistinguishable from that of the strongly balanced sample.

[9] This is somewhat in contrast to the role of the extensive margin for employment dynamics as documented by Fort, Pierce, and Schott (2018): While entry and exit (of establishments within firms or firms altogether) may account for 88% of employment changes in U.S. manufacturing, labor shares of entrants and exiting establishments are not different enough from that of incumbents, and the value added they account for is not large enough for them to impact the manufacturing labor share decline.

in the following decades, a scenario we label the "superstar economy." Both are examples of the composition effects we discussed in Section 3.1, and compatible with the two facts highlighted under Finding 1: (1) a relatively stable median labor share and (2) a larger portion of manufacturing output produced at the bottom of the labor share distribution.[10]

To assess which of these two scenarios, if any, is quantitatively relevant, we construct two counterfactual aggregate labor shares. First, we compute the manufacturing labor share by keeping an establishment's value-added weight ω_{it} equal to its 1982 value[11], while allowing its labor share λ_{it} to evolve over time as its does in the data:

$$\lambda_t^{BIG} = \sum_i \lambda_{it}\omega_{i1982}.$$

In the second exercise, we do the reverse: the labor shares λ_{it} are kept constant to their values in 1982, while we use the historical time series of the weights ω_{it}. Specifically, this second counterfactual labor share is defined as

$$\lambda_t^{STAR} = \sum_i \lambda_{i1982}\omega_{it}.$$

In order to conduct this analysis, we focus on a strongly balanced sample of establishments that are permanently active from 1982 to 2012. We display the actual and counterfactual labor shares in Figure 5.

Figure 5: The limited role of establishments with an initially small size or low labor share

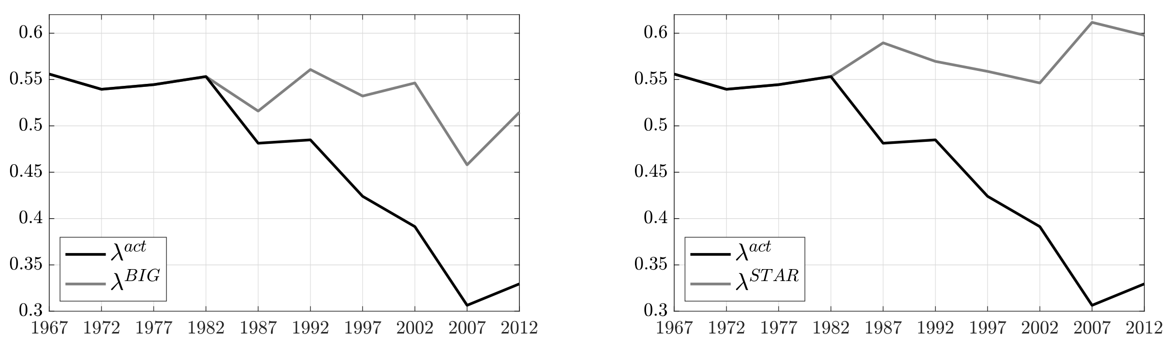

The main takeaway from Figure 5 is clear: none of the two counterfactuals is able to reproduce the actual drop in the labor share of the U.S. manufacturing sector, λ_t^{act}. In the "big-player economy" (λ_t^{BIG}), the counterfactual labor share is actually a little *higher* in 2012 than it was in 1982, while it dropped by about 18 ppt in the data over the same time period. In other words, the fall in the aggregate labor share does not appear to be driven by a divergence in the relative labor shares of large versus small establishments.

[10]We thank an anonymous referee for suggesting these scenarios.

[11]As a robustness check, we also consider their average values around that Census: $\overline{\omega}_{i1982} = E_i[\omega_{i\tau}], \tau = 1977, 1982, 1987$. In the next exercise, we do the analogue with labor shares.

If we instead freeze the labor shares and focus on the contribution of the value added weights (the "superstar economy" λ_t^{STAR}), we obtain an even more noticeable *rise* in the counterfactual labor share: from 56 ppt in 1982, it increases to 62 ppt by 2012, in contrast to a drop to 38 ppt in the actual data. The takeaway: manufacturing establishments with an initially low labor share (the *superstars*) did not experience, on average, higher growth in value added between 1982 and 2012. As we saw in Section 3.2, this could have happened, for example, if they had been more prone to experience positive TFP shocks over this period.

4.3.3 The role of the joint dynamics

Finding 1 made clear that the decline in the manufacturing labor share was due to a strong decline in the covariance between labor and market shares at the establishment level. The conclusion from the two exercises in the previous section is that *neither* market share *nor* labor share dynamics at the establishment level can, *on their own and separately*, explain this development. Instead, the *joint* dynamics of labor shares and size at the micro level are key to our understanding of the aggregate trends in $Cov(\lambda_{it}, \omega_{it})$ and the labor share. To highlight their importance, we next perform a shift-share decomposition among incumbent establishments, in the spirit of e.g. Baily, Hulten, and Campbell (1992):

$$\begin{aligned}
\text{Total: } -18 \text{ ppt} \qquad \lambda_{2012} - \lambda_{1982} &= \sum_i \omega_{i1982}(\lambda_{i2012} - \lambda_{i1982}) && \text{Shift: } +1.5 \text{ ppt} \\
&+ \sum_i (\omega_{i2012} - \omega_{i1982})\lambda_{i1982} && \text{Share: } +4.5 \text{ ppt} \qquad (10) \\
&+ \sum_i (\omega_{i2012} - \omega_{i1982})(\lambda_{i2012} - \lambda_{i1982}) && \text{Joint: } -22.8 \text{ ppt.}
\end{aligned}$$

The first term in equation (10) is akin to the earlier "big-player" scenario, in which the value-added weights are kept constant to their values in the 1982-Census, ω_{i1982}. The second term instead freezes the micro-level labor shares at λ_{i1982} while allowing weights to fluctuate, similar to the "superstar economy" counterfactual. The last term corresponds to the interaction between the changes in the labor share and changes in the value added share.

The numbers that accompany equation (10) help us quantify our earlier conclusions. Specifically, the first two terms contribute *positively* to the evolution of the aggregate labor share, across all sample periods considered: For example, the contribution of the change in the labor shares holding weights constant is +1.5 ppt per decade over the 1967-2012 period, while that of the change in the weights alone is +4.5 ppt per decade.

In order to explain the dramatic decline in the manufacturing labor share, it must therefore be that the contribution of the interaction term is strongly negative. Indeed, over the whole sample period, its contribution is −22.8 ppt per decade on average. This term explains *more than* the 18 ppt overall decline in the labor share.

The main takeaway from the findings in this section is that the decline in the aggregate labor share is not simply due to a composition effect whereas superstar establishments take over the market. Instead, it appears to be driven by units whose labor shares fall *at the same time* as they grow in size. These joint dynamics led to a strong reallocation of economic activity towards the left tail of the labor share distribution, and caused $Cov(\lambda_{it}, \omega_{it})$ to turn negative. They also implicitly suggest a polarization of labor shares across establishments, rationalizing the fattening of the tails of the (unweighted) labor share distribution that we described under Finding 1.

What could be behind these joint dynamics? The conceptual framework of Section 3 provides a few candidates. For example, establishments facing non-isoelastic demand schedules may have experienced strong positive demand or TFP shocks. As we discussed, this would be expected to generate higher markups, leading to a fall in those establishments' labor shares λ_{it} and a rise in their economic weights ω_{it}.

Yet, distinguishing between the various scenarios that we analyzed in Section 3 ultimately requires a deeper analysis of the micro-level dynamics of labor shares and value added shares. This is what we turn our attention towards for the rest of this paper. In the next section, we start by showing that the typical superstar low-labor-share establishment has a surprisingly transient labor share. As such, these establishments are more akin to "shooting stars" than superstars.

4.4 Finding 3: Low labor shares are transient

So far, our analysis of the labor share decline has been static in nature: we have focused on cross-sectional snapshots of the data 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. This is not innocuous: the evidence from Figure 5 and the shift-share decomposition in Equation (10) made it clear that these dynamics are central to understanding labor share trends at higher levels of aggregation. In turn, they can teach us about the forces and factors that lie behind the decline in the labor share.

4.4.1 How persistent are establishment labor shares?

As a starting point for our analysis of micro-level dynamics, we study the persistence of manufacturing establishments' labor shares. Our first exercise is to regress a unit's labor share on its past level as well as a number of controls:

$$\lambda_{it} = (1-\rho)\left(\overline{\lambda} + \alpha_i\right) + \rho\lambda_{it-5} + X_{it} + \varepsilon_{it}. \tag{11}$$

The coefficient α_i corresponds to the permanent deviation of an establishment's labor share from the average level, $\overline{\lambda}$; ρ is the persistence of a labor share's deviation from its long-run average; and X_{it} is a vector of industry, state and year controls. We estimate equation (11) on the full panel over the 1967-2012 period using the consistent estimator by Arellano and Bond (1991) to address potential issues due to the serial correlation between λ_{it-5} and the error term. We report results in Table 1, for both an unweighted specification and regressions with value added weights at the plant-year level.

Table 1: AB estimates of labor share persistence

	(1)	(2)
Data	CMF	CMF
$\bar{\lambda}$	0.784	
ρ	0.0817	0.1756
	(0.0022)	(0.0023)
Estimation	AB	AB
Weights	No	Yes

Note: Data are from the Census of Manufactures and cover the 1967-2012 sample. Results are from a panel regression with fixed effects using the consistent estimator by Arellano and Bond (1991). Controls include a full set of industry, state and year dummies.

To ease the interpretation of the results, let us rewrite equation (11) in gap form:

$$\lambda_{it} - \bar{\lambda} = (1-\rho)\alpha_i + \rho(\lambda_{it-5} - \bar{\lambda}) + X_{it} + \varepsilon_{it} \tag{12}$$

where $\bar{\lambda}$ is the relevant industry/region/time long-term labor share for each plant.

The low coefficients on the autoregressive term suggest that labor shares at the establishment level are very transient. The coefficient estimates for ρ imply that only between 8.2% and 17.6% of a plant's current labor share deviation from its long-run level remains five years later. For example, consider an establishment with $\alpha_i = 0$ that has a labor share today of 40 ppt, while the average labor share in its industry and region that year is 60 ppt. Then, the weighted annual estimate in the second column implies that its labor share would be expected to be 56.5 ppt five years later, assuming that the industry/region labor share has remained constant over time.

This evidence indicates that micro-level labor share dynamics are not only central to our understanding of aggregate trends, but also display a surprising degree of transience. Next, we study dynamics at the bottom end of the labor share distribution.

4.4.2 Defining low-labor share (LL) establishments

We saw in Section that the bottom of the labor share distribution experienced a dramatic rise of its economic importance between 1967 and 2012. These are the establishments that we will be focusing on for much of the following analysis. Specifically, we define low-labor (LL) establishments as those in the lowest quintile of the labor share distribution in a given year and 3-digit NAICS industry. The quintiles are industry-specific due to the wide range of average labor shares across industries.

To highlight the role of LL establishments in shaping aggregate dynamics, we start by recomputing the manufacturing labor share without them. If reallocation towards lower labor share establishments 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 LL establishments

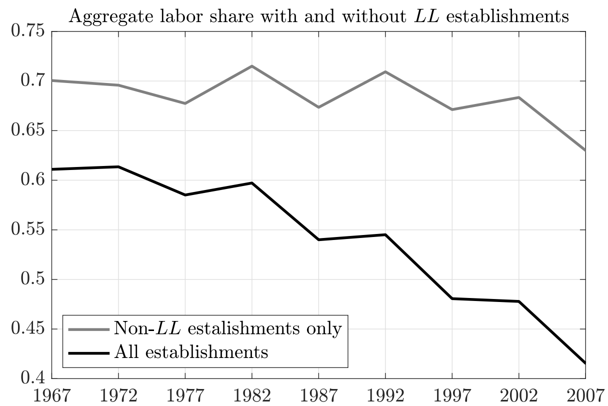

Note: The figure plots the manufacturing labor shares computed on the full panel (solid black line) against that computed for the panel after dropping the set of LL establishments (solid grey line). It shows that Non-LL establishments do not contribute to the decline of the manufacturing labor share.

The labor shares including and excluding LL establishments are shown in Figure 6. Two aspects stand out: First and not surprisingly, the level of the manufacturing labor share without the bottom quintile of the distribution is much higher at about 0.75. Second, and more importantly, it does not exhibit any significant decline: while the actual manufacturing labor share starts to fall in the 1980s, the counterfactual manufacturing labor share without LL establishments fluctuates around its time-series mean, with no clear trend. In other words, while reallocation among Non-LL establishments may be taking place, it does not contribute meaningfully to the empirically observed manufacturing labor share decline. This indicates that analyzing the nature of LL establishments is key to understanding the forces behind the manufacturing labor share decline.

Now that we have defined LL establishments, we aim to characterize their properties. More specifically, we show that they are akin to economic "shooting stars" – their labor share and economic importance shine brightly only for a limited time before reverting back to normal. Yet, as will become clearer later, they remain quantitatively relevant for manufacturing labor share trends despite their temporary nature.

4.4.3 Markov transitional dynamics

We start by documenting the transition dynamics of LL and Non-LL establishments with the help of a Markov transition matrix, displayed in Table 2. That is, we ask a simple question: conditional on an establishment's labor share at time t, what is the probability that it has LL status at time $t+5$? If establishment-level labor shares were highly persistent and LL establishments retain there status over time, this probability should be equal to 100%. At 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 LL in terms of the lowest quintile).

Table 2: Transition probabilities of LL status

Panel A. Unweighted transitional dynamics		
	Non-LL_{t+5}	LL_{t+5}
Non-LL_t	0.854	0.146
LL_t	0.583	0.417
Panel B. Weighted transitional dynamics		
	Non-LL_{t+5}	LL_{t+5}
Non-LL_t	0.922	0.078
LL_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 LL status when quintiles are unweighted, Panel B. displays the share of manufacturing value added accounted for by the LL establishments when defined by VA-weighted quintiles.

Table 2 shows that over our sample period, the probability that an establishment retains LL status from Census year to Census year (a 5-year window) is only 41.7%. While this is higher than if LL 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.

One may be concerned that the results in Table 2 are mostly driven by small, economically insignificant establishments. For this reason, we also consider Markov transition matrices of quintiles weighted by economic activity and confirm the transient dynamics of LL 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 remains unchanged.

4.4.4 Labor share dynamics of LL establishments: Shooting stars, not superstars

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

$$\Delta\lambda_{it} \equiv \lambda_{it} - \lambda_{it-5} = c_1 + \beta_{-5}\mathbb{I}\{LL_{it}\} + \gamma_1 X_{it} + \varepsilon_{1it} \quad (13)$$

$$\Delta\lambda_{it+5} \equiv \lambda_{it+5} - \lambda_{it} = c_2 + \beta_{+5}\mathbb{I}\{LL_{it}\} + \gamma_2 X_{it} + \varepsilon_{2it} \quad (14)$$

While the *level* of the labor share of LL 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 relative *dynamics* from the estimates of the coefficients β_{-5} and β_{+5} in Equations (13) and (14). That is, we study how the labor share dynamics of LL establishments differ from those of Non-LL establishments over a ten-year window around the reference period.

Note that we do not require that LL establishments in period t were also LL establishments in $t-5$ and will be in $t+5$; an establishment could well have LL status for a single year. The vector X_{it} contains industry, region and year dummies as controls. We estimate Equations (13) and (14) with and without value-added weights to account for the fact that larger establishments are likely to have less volatile labor shares. Results are displayed in Table 3.

Table 3: The dynamics of LL 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 (13) and (14) on the full Census panel. "VA weights" correspond to the share of establishment i's value added in overall manufacturing 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 LL 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 LL and Non-LL 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. This confirms that the low overall persistence of labor shares we documented earlier is also present at the low end of the distribution.

The coefficient estimate from the unweighted regression in column (I) implies that relative to the previous Census year, an establishment that has LL 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-LL 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 LL in year t will be 27.2 ppt *higher* than that of Non-LL establishments. Columns (III) and (IV) in Table 3 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). 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 7. It is striking to see that while the relative labor share change of time-t LL establishments is very negative between $t-5$ and t, by $t+5$ these LL establishments appear to be no more different than their Non-LL peers than they were 10 years earlier. All in all, our analysis appears to show that the average LL 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. Before we turn our attention to understanding the sources and implications of this volatility, we show evidence that it is not an artefact of pervasive measurement error.

Figure 7: The temporary fall and rise of labor shares of LL establishments

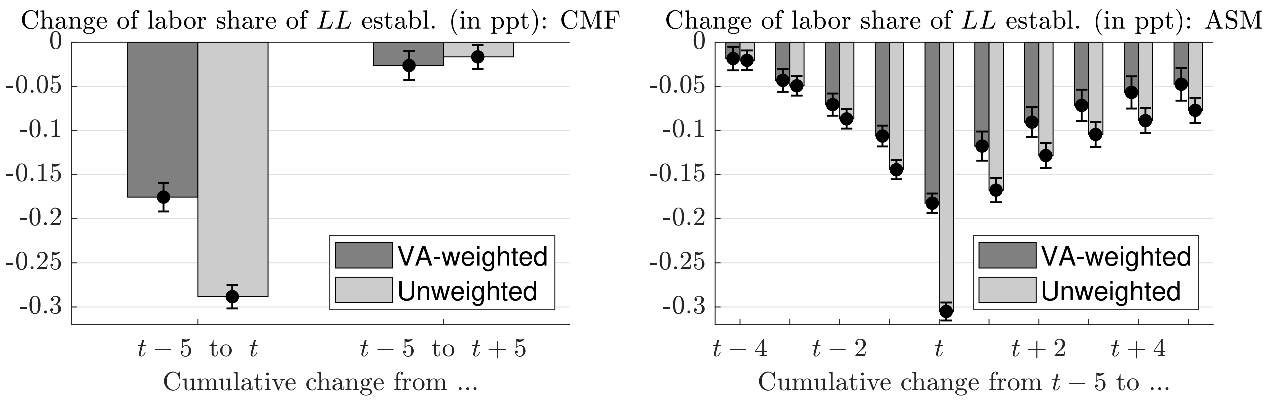

Note: Left panel: Cumulative evolution of the labor share of the average LL establishment relative to their peers in the Census panel before ($t-5$ to t) and after (t to $t+5$) the year it is in LL status. Unweighted dynamics in dark grey, value added-weighted dynamics in light grey; whiskers denotes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

Labor share dynamics and measurement error. One potential concern is that the low persistence of the labor share is driven by widespread measurement error. Under this scenario, LL_t establishments would simply be establishments that experienced large (negative) 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 7. Using only data every five years would make our analysis vulnerable to measurement errors in just that single year. While this may be a concern especially for small establishments, measurement error for large establishments whose labor share is low is much less likely as Census pays a lot of attention to large producers that matter greatly for their aggregate tabulations.

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 LL status were merely driven by idiosyncratic measurement error, we would expect establishments that are LL establishments to look on average like their non-LL 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 (13) and (14) 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 7: 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-LL establishments. All in all, our evidence appears to indicate that the transient nature of LL status is not merely an artifact of transient measurement error.[12]

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 LL firms analogously to our establishment-level definition. The results are displayed in Appendix B.4 and are largely similar to the ones we discussed in this section.

4.5 Finding 4: Labor share dynamics are driven by value added, not wages or employment

In the previous sections, we documented that low-labor share establishments have accounted for an increasingly large portion of manufacturing value added and that LL status is remarkably transient. Going back to the conceptual framework of Section 3, Findings 1 to 3 have implications about the most likely shocks behind the decline in the manufacturing labor share: they must (1) generate a negative correlation between labor share and value added dynamics at the establishment level and (2) be temporary. All three types of shocks we discussed in Section 3.2 – demand, technology or monopsony – are consistent with the previous evidence. To discriminate between them, we now turn our attention to the components of the labor share.

4.5.1 Wages and labor productivity across establishments

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

$$\log \lambda_{it} = \log W_{it} - \log ARPL_{it} \tag{15}$$

where W_{it} is the wage of the average employee and $ARPL_{it} = P_{it}Y_{it}/L_{it}$ denotes the average revenue product of labor.[13]

To ensure that our results are not driven by systematic differences across industries, regions or time as well as to make them more readily interpretable – wages and value added per worker

[12] 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 LL establishments using this alternative sales measure are very similar to our benchmark presented in Figure 7. 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.

[13] It is important to notice that both W_{it} and $ARPL_{it}$ are nominal variables; the latter compounding 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).

are nominal variables –, we study an establishment's wage, value added per worker relative to that 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.[14] The relative wage, $\widetilde{w}_{it}$, and labor productivity, $\widetilde{p_{it}y_{it}/l_{it}}$, are then defined in logs as follows:

$$\widetilde{x}_{it} \equiv \log X_{it} - \overline{\log X_{-i,t}} \quad \text{ where } \overline{\log X_{-i,t}} \equiv \sum_{j\neq i} \frac{P_{jt}Y_{jt}}{\sum_{j\neq i} P_{jt}Y_{jt}} \log X_{jt} \text{ and } X_{it} = W_{it},\ \frac{P_{it}Y_{it}}{L_{it}}. \tag{16}$$

where we omitted the industry and region subscripts for expositional purposes. The measures $\widetilde{x}_{it}$ are by definition 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.

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

$$\widetilde{x}_{it} = f(\lambda_{it}) + \varepsilon_{it}, \qquad\qquad \widetilde{x}_{it} = \widetilde{w}_{it},\ \widetilde{p_{it}y_{it}/l_{it}}, \tag{17}$$

where $\widetilde{x}_{it}$ is either establishment i's relative wage, $\widetilde{w}_{it}$, or labor productivity, $\widetilde{p_{it}y_{it}/l_{it}}$. The function $f(\cdot)$ is the object of interest: It indicates whether low-labor-share establishments pay lower wages than their peers and/or experience higher labor productivity. To ensure that we measure economically-relevant relationships, each observation is weighted by the establishment's share in manufacturing 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 $\widetilde{py/l}$ are scaled within each industry, year and region, we ensure that our findings are not driven by these factors.

The results of the three non-parametric regressions are displayed in the left panel of Figure 9. They paint two clear and striking pictures. First, relative wages are nearly orthogonal to the labor share: LL establishments do not on average pay their workers more or less than their peers.[15] 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: $\widetilde{py/l}$ 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 ($\widetilde{py/l} = 0$) at a labor share of $\lambda = 0.46$. These differences are large. For example, establishments with a labor share

[14] 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.

[15] 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.

of $\lambda = 0.1$ experience a relative labor productivity of $\widetilde{py/l} = 1.04$, meaning that they produce $\exp(1.04) \approx 2.83$ times more value added per worker than the average establishment in the same industry, region and year. At the other end of the spectrum, establishments with a labor share of unity exhibit $\widetilde{py/l} = -0.61$, which means that they produce only a bit above half the value added per worker ($\exp(-0.61) \approx 0.54$) of their peers. LL establishments have an average relative labor productivity of 0.596 compared to -0.428 for Non-LL establishments; the average LL establishment thus produces about 2.8 times more value added per worker than the typical Non-LL establishment. Yet, in terms of relative wages, they do not differ at all.

Figure 8: Labor productivity dominates cross-sectional differences of labor shares.

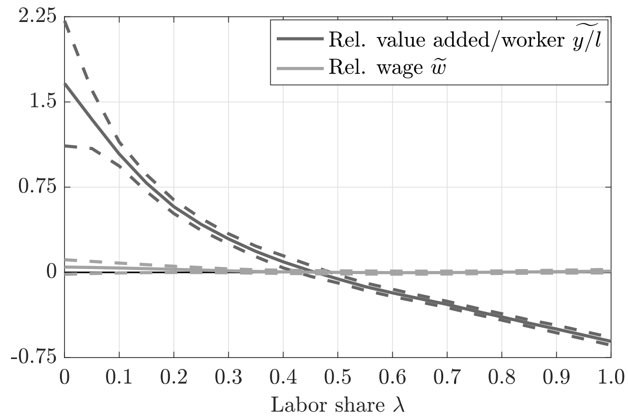

Note: The figure displays the cross-sectional differences in relative value added per worker $\widetilde{py/l}$ and the relative wage $\widetilde{w}$ against the labor share. All relative measures denote log-point differences vis-à-vis their peers as defined in Equation (16). Dashed lines denote 95% error bands.

4.5.2 Dynamic evidence

Next, we investigate which components are behind the V-shaped pattern of the LL establishments' labor shares that we documented earlier. Starting from the relationship in growth rates:

$$\begin{aligned}\Delta \log \lambda_{it}^{LL} &= \Delta \log W_{it}^{LL} + \Delta \log L_{it}^{LL} - \Delta \log Y_{it}^{LL} \\ &= \Delta w_{it}^{LL} + \Delta l_{it}^{LL} - \Delta y_{it}^{LL},\end{aligned}$$

we can apply the same regression strategy that was used to produce Figure 7 in order to decompose the change in the labor share of LL establishments at time t into the responses of wages, employment and value added.

Figure 9 shows that labor productivity is not only behind the cross-sectional differences along the labor share dimension, but also the main source of time-series variation in the labor share of LL establishments. Recall that over the five years prior, the average LL establishment saw its labor share shrink by 18 ppt compared to the typical Non-LL establishment. The evidence in Figure 9 clearly indicates that this is almost entirely driven by differential value added growth, with no

Figure 9: Value added dynamics dominate labor share dynamics of LL establishments.

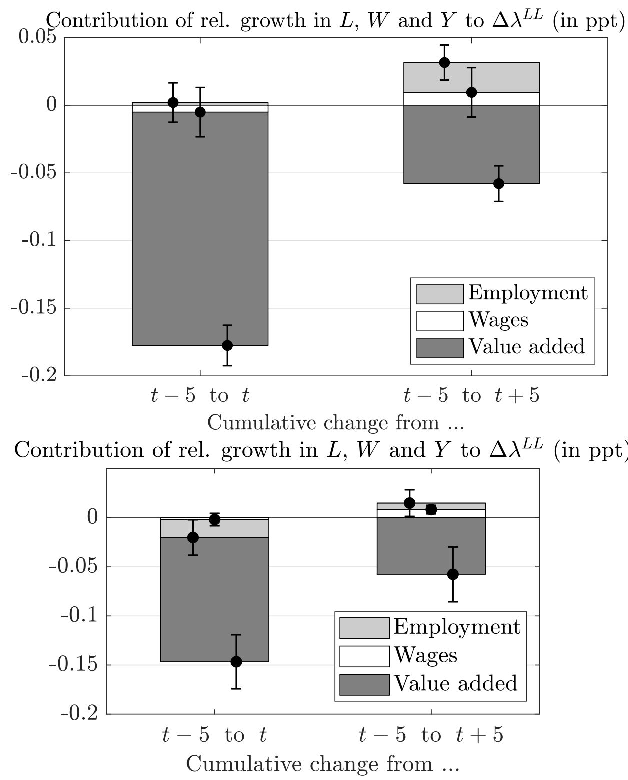

Note: This figure displays the dynamic contributions of the changes in wages, employment and value added for labor share dynamics of the average LL establishment relative to its peers. The first bars display their contributions before ($t-5$ to t), the second bars their cumulative contributions until after ($t-5$ to $t+5$) the year an establishment is in LL status. Whiskers denote 95% error bands.

statistically significant contribution from wages or employment. Incorporating the five following years, we see that the V-shape pattern of the labor share between $t-5$ and $t+5$ is mainly a result of the reversal of the initial jump in value added of LL establishments. This retreat of value added growth accounts for 12 ppts of the 15-ppt rebound in the average LL labor share. In other words, there is little contribution coming from a delayed response of employment growth (contribution of +2 ppt) or wages (+1 ppt).[16] The contribution of value added dynamcis to labor share dynamics becomes even more dominant over time with less and less of a role of wage and employment dynamics.

In sum, we find no evidence that wages are behind labor share differences across establishments or driving establishment-level labor share dynamics. Going back to the framework of Section 3.2, it appears therefore unlikely that increased monopsony power in the labor market is behind the fall in the manufacturing labor share.

4.6 Finding 5: Low labor shares stem mostly from a "product price premium"

The previous sections revealed a number of puzzling findings. Not only did we show that the manufacturing labor share declined due to a dramatic reallocation of value added towards LL 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 shooting stars? Finding 4 gave us a hint about the cause: cross-sectional and dynamic differences between LL and Non-LL appear to be driven by *nominal* value added per worker. This leaves two candidate forces driving the manufacturing labor share decline: nominal price dynamics and real labor productivity. Next, we attempt to disentangle these two opposite forces and provide evidence that demand-side factors rather than technology appear to be a key driver of micro-level labor share patterns.

4.6.1 Measuring prices

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`). In addition, it 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.[17] This is a product-coding system

[16] We also studied the dynamics of capital intensity and intermediates. We find little evidence that they are different for LL establishments relative to their peers.

[17] Foster, Haltiwanger, and Syverson (2008) define products at the 7-digit SIC code level while Bernard, Redding,

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.5).[18]

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 manufacturing goods 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 earlier found to be the key driver of labor shares in the cross section 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 (16):

$$\widetilde{p_{it}q_{it}/l_{it}} \equiv \log\left(P_{it}Q_{it}/L_{it}\right) - \overline{\log(P_{-i,t}Q_{-i,t}/L_{-i,t})} \tag{18}$$

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

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.6 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.

4.6.2 Product prices across establishments

In order to make prices comparable across establishments, we adopt the treatment of nominal wages and labor productivity in Section 4.5 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

and Schott (2010, 2011) aggregate product sales to the 5-digit SIC level of products; both definitions are coarser than ours.

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

10-digit NAICS product ℓ:

$$\widetilde{p}_{i\ell t} \equiv \log P_{i\ell t} - \overline{\log P_{-i,\ell t}} \qquad \text{where } \overline{\log P_{-i,\ell t}} \equiv \sum_{j\neq i} \frac{P_{j\ell t} Q_{j\ell t}}{\sum_{j\neq i} P_{j\ell t} Q_{j\ell t}} \log P_{j\ell t} \tag{19}$$

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 its product lines. This measure represents the mean log-point difference between an establishment's output prices and those of its peers.[19]

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 10.

In addition, we run regressions analogue to Equations (13) and (14) 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 10.

We can see from the left panel that while relative sales per worker are driven by both price and physical labor productivity differences, price are 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 a labor share of 50% and more. Focusing on LL establishments, their average relative price is 0.15 compared to -0.041 for Non-LL establishments. This difference implies a product price premium for LL establishments of roughly $\exp(0.15 + 0.041) \approx 21\%$. This contributes a fair amount to the relative sales per worker that LL establishments generate to their peers of $\exp(0.430 + 0.096) \approx 69\%$.

[19] 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.

Figure 10: Relative labor productivity and relative prices

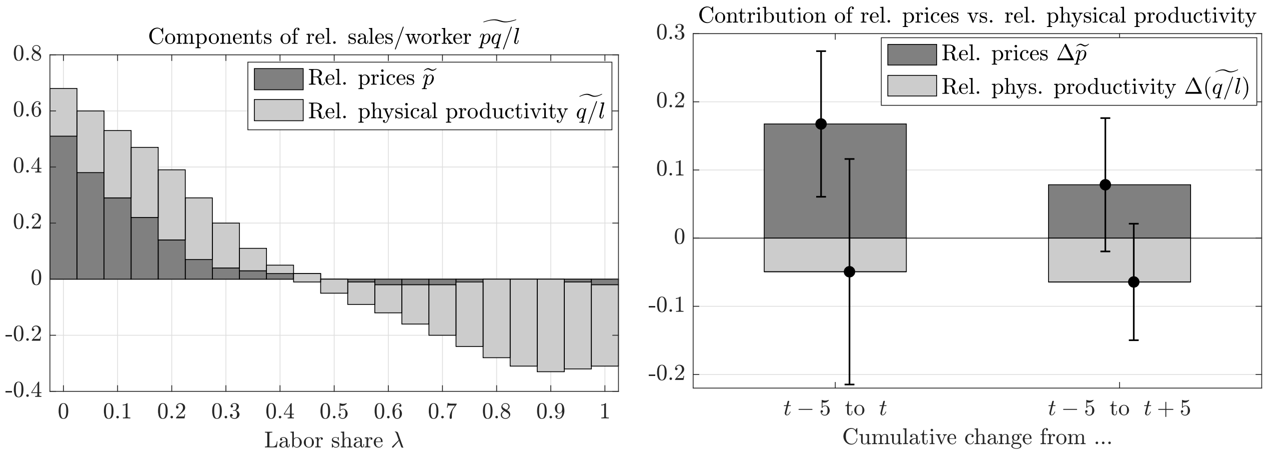

Note: Left panel displays the cross-sectional differences in relative prices $\widetilde{p}$ (dark grey bars) and relative physical labor productivity $\widetilde{q/l}$ (light grey bars) against the labor share; $\widetilde{p}$ defined in Equation (19) and $\widetilde{q/l}$ is defined as the ratio of $\widetilde{pq/l}$ (defined in Equation (18)) 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 LL establishment relative to their peers. The first bars display their cumulative contributions before ($t-5$ to t) and after ($t-5$ to $t+5$) the year an establishment is in LL status. Whiskers denote 95% error bands.

4.6.3 Dynamic evidence

The right panel of Figure 10 displays the dynamic analysis. Analogous to the regression in Equations (13) and (14), we estimate how the price and sales per worker dynamics of time-t LL establishments differ from those of Non-LL establishments. Again, we infer the physical productivity growth of LL establishments as a residual. The results of the dynamic analysis are even starker than those from the cross-section: compared to their Non-LL peers, the relative prices of LL 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 LL at time t relative to their peers. While the LL-vs-Non-LL 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 -5% and are statistically insignificant.

All in all, the evidence in this section indicates that LL status stems from an exceptionally high *revenue* productivity[20], generated through a capacity to extract a lot of revenue out of their workforce by charging their customers higher prices.

The findings in this section provide important insights that help us discriminate between the potential theories behind the dramatic decline in the manufacturing labor share. From the framework of Section 3.2, we know both demand- and technology-based theories could be compatible with

[20]In Appendix B.4 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.

Findings 1 to 4: transient preference or TFP shocks combined with non-isoelastic demand schedules can explain the joint dynamics of labor and market shares at the establishment level; the reallocation of economic activity towards low-labor-share establishments; and the dominant role played by value added in contrast to employment or wages. Yet, the fact that relative prices and labor shares co-move negatively represents strong evidence in favor of demand shocks: under technology shocks, we would expect relative prices to *fall* alongside labor shares. Importantly, the outsized contribution of relative prices to the levels and dynamics of value added for LL establishments suggests that demand shocks are first order to explain labor share dynamics.

5 Discussion: Why did the manufacturing labor share decline?

The previous sections documented evidence pointing at the prominent role of demand shocks in shaping relative product prices and labor share dynamics. How do some establishments manage to extract high prices compared to their peers? And how did the dynamics of the labor share components change over time? In this section, we explore a potential avenue and provide some supporting evidence.

5.1 Finding 6: Relationship between advertising activity and labor share

The marketing literature has documented that firms use advertising to gain market power through multiple channels, e.g. perceived differentiation, customer retention or brand recognition (REFERENCES). In theory, marketing activity can help the firm not only engineer a "preference shock" to lift demand for its products by spreading information, but also build a loyal customer base that insulates its demand from price increases. Next, we show that advertising intensity is strongly related to labor share, hinting at its role in driving the dynamics we documented earlier.

5.1.1 Cross-sectional evidence

Manufacturing establishments are asked to report information about their advertising expenditures since 1997. One has to be mindful that advertising expenses at the head-quarter level are missing in our analysis. As this likely weakens the relationship between advertising expenses and labor shares, we view our results as a lower bound on the true effect of advertising.

We compute establishment i's advertising expenses per employee in year t, ca_{it}/l_{it}, and scale that number analogously to the other nominal variables as illustrated in Equation (16). The resulting variable, $\widetilde{ca_{it}/l_{it}}$, denotes the log point difference of an establishment's advertising expenditures per employee relative to that of their peers in the same industry, state and year. As in the previous sections, we non-parametrically regress that variable on the labor share and plot the estimates in Figure 11.

The figure reveals a clear pattern: low labor share establishments spend much more on advertising compared to their peers than high labor share establishments do. Our estimates indicate that a typical unit with a labor share of 0.1 spends about $\exp(0.51) = 1.67$ times as much on advertising

Figure 11: Cost of advertising per employee and labor shares

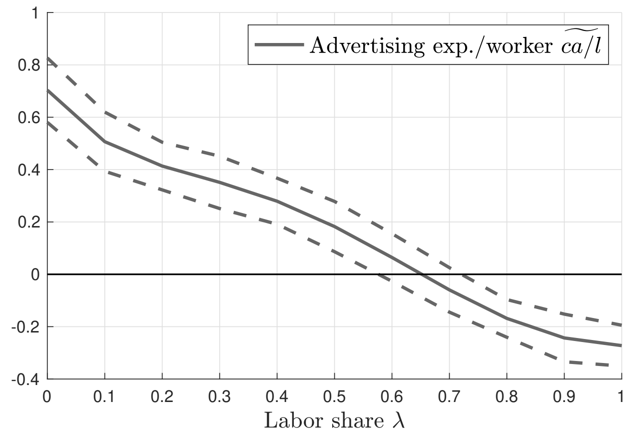

Note: Data come from ...

as the average plant in its sector and region, while high labor share establishments spend about 0.25% less.

5.1.2 Dynamic evidence

The availability of multiple cross-sections allows us to study next the dynamic effects of advertising expenditures on future labor shares. Specifically, we run the following regression:

$$\mathbb{I}\{LL_{it}\} = c + \beta_1 \widetilde{ca/l}_{it-5} + \beta_2 t + \beta_3 \widetilde{ca/l}_{it-5} \times t + \gamma X_{it} + \varepsilon_{it} \tag{20}$$

Table 4: The increasing impact of advertising intensity on labor shares

Variable	(I)	(II)
β_1	0.0216***	0.0218***
	(0.0030)	(0.0061)
β_3	0.0028***	0.0008
	(0.0010)	(0.0027)
R^2	0.111	0.096
Weights	none	VA weights

Note: Pooled OLS regression of Equation (20) in years 1997-2012 on the full Census panel. For details, see notes to Table 3.

The unweighted estimate of β_1 is positive, significant and implies that a one standard deviation in advertising intensity, $\widetilde{ca/l}$, implies that the probability of becoming a LL establishment five years later increases by one tenth. This suggests that an establishment's product prices benefit from previous advertising expenditures.

In addition to that unconditional statement, the coefficient estimate of β_3 is positive. This means that the sensitivity of becoming an LL establishment to advertising expenditures has increased over time. Compared to 1997, the sensitivity is 40% higher in 2012. This increased effectiveness of advertising expenditures sounds plausible given advertising benefitted from a wider reach and better targeting in recent years. Online media have made it possible to address customers throughout the country (or even the world) compared to the traditional high way billboard or ad in the local newspaper. Finally, column (II) displays weighted estimates. With the exception of the significance of β_3 the results here are similar to the unweighted ones.

Not only has the effectiveness of advertising expenses appeared to have risen over time, but there is evidence that marketing intensity has also been growing: according to Statisa, the resources spent on advertising increased at a rate of 7.5% per year between 2004 and 2018, easily outpacing GDP growth.

5.2 Finding 7: The V-shaped pattern gets deeper over time

If the demand factors behind the labor share dynamics are partly driven by marketing activity, one would expect that the rising effectiveness and intensity of advertising over time should amplify their importance. Next, we investigate the evolution over time of the V-shaped pattern that we documented earlier.

For each Census year t we compute the change in the labor share between $t-5$ and t for time-t LL ($\Delta\lambda_t^{LL}$) and Non-LL ($\Delta\lambda_t^{\text{Non-}LL}$) establishments, and plot the difference. Both unweighted and value-added-weighted estimates are shown in Figure 12. In both cases, we find that the labor share dynamics of LL and Non-LL get increasingly different over time, 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).

5.3 Finding 8: Employment has become more disconnected from value added

In Finding 4, we showed that nominal labor productivity was central to understanding the labor share response of LL 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 manufacturing labor share (2000s) relative to the early part of the sample when the labor share was more stable (1970s).

First, we 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 (17), 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."

Figure 12: Labor share change of LL versus Non-LL establishments over time

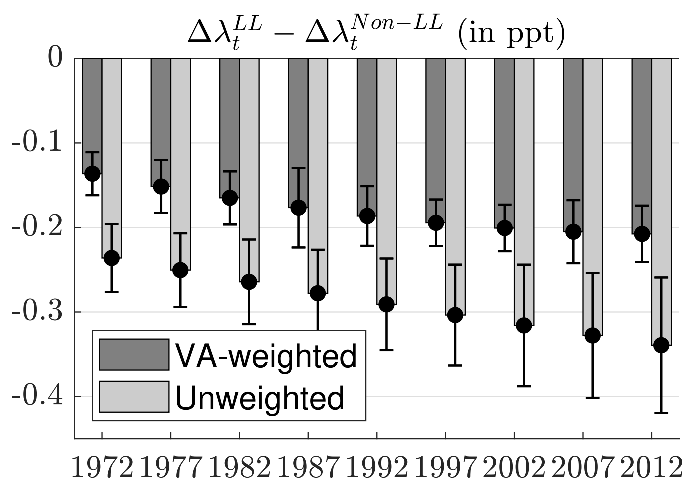

Note: This figure displays the difference in labor share dynamics between LL and Non-LL establishments (corresponding to the $t-5$ to t bars in the left panel of Fig. 7) by year. It shows that LL establishments look increasingly different from their peers starting in 1987.

Figure 13: 1970s vs. 2000s.

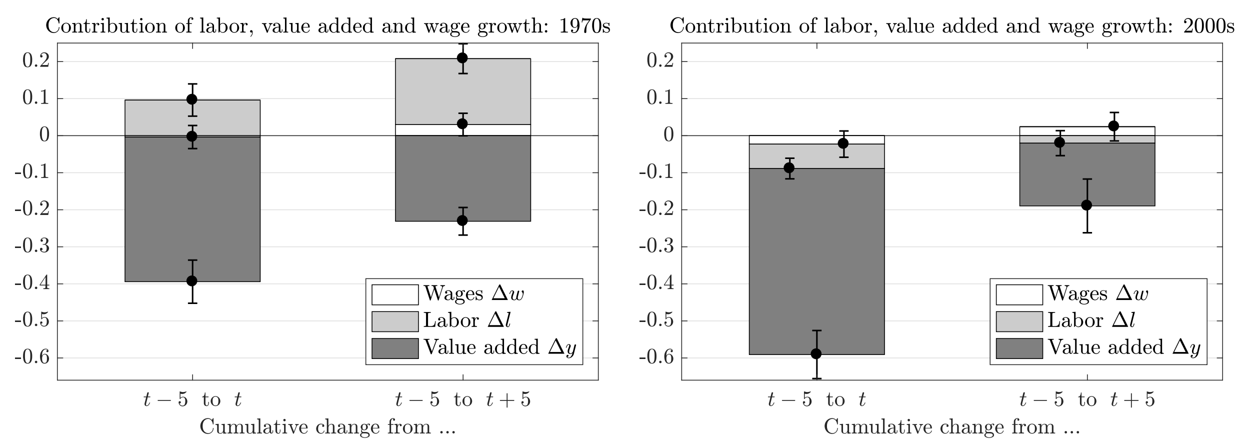

The results are reported in the top row of Figure 13. A few things are worth noting. First, the link between wages and labor share was always weak, both economically and statistically. Second, while the inverse relationship between labor share and productivity was always strong, it became even more pronounced over time; the nonparametric estimate is particularly steep in the 2000s (top-right panel) for the establishments with the lowest labor share. Third, it is worth noting that the labor productivity curve crosses the x-axis at a lower labor share in the latter part of the sample. This 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 (13) and (14), but this time applied separately to wages, employment and value added and study how LL establishments differ from their peers in the dynamics of these variables:

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

The middle panels of Figure 13 display the contribution of these three components to the labor share growth rate of LL establishments relative to non-LL establishments. On the left, we can see 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 LL establishment's labor productivity growth was 40 ppt higher than that of non-LL 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 LL 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 LL establishments whose hiring seems to respond to the strong value added growth, but with a delay. Ultimately, while the value added of LL 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 LL 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 LL establishments relative to their Non-LL 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). It also gives credence to our hypothesis that the dynamics of value added are driven by markup variations due to advertising activity.

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 manufacturing 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 superstar, LL, 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 LL establishments (1) have highly transient labor shares; (2) are more and more likely to have been large before earning their superstar 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 manufacturing 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 7.)

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 pages 7 and 11.)

Manuel Arellano and Stephen Bond. Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations. *Review of Economic Studies*, 58(2):277–297, April 1991. (Cited on pages 21 and 22.)

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 pages 8 and 11.)

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 8, 11, 17, and 48.)

José Azar, Ioanna Marinescu, and Marshall I. Steinbaum. Labor market concentration. *NBER Working Paper No. 24147*, 2017. (Cited on page 12.)

Martin N. Baily, Charles Hulten, and David Campbell. Productivity dynamics in manufacturing plants. *Brookings Papers on Economic Activity*, 1992:187–267, 1992. (Cited on page 20.)

David Rezza Baqaee and Emmanuel Farhi. Productivity and misallocation in general equilibrium. *Quarterly Journal of Economics*, forthcoming. (Cited on pages 8 and 11.)

Simcha Barkai. Declining labor and capital shares. *Working Paper*, 2017. (Cited on pages 8 and 11.)

David Berger, Kyle Herkenhoff, and Simon Mongey. Labor market power. *NBER WP No. 25719*, March 2019. (Cited on page 12.)

Daniel Berkowitz, Hong May, and Shuichiro Nishioka. Does capital-labor substitution or do institutions explain declining labor shares? *Working Paper*, 2017. (Cited on page 17.)

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 31.)

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 32.)

Olivier Blanchard and Francesco Giavazzi. Macroeconomic effects of regulation and deregulation in goods and labor markets. *Quarterly Journal of Economics*, 118(3):879–907, August 2003. (Cited on page 12.)

Petri Böckerman and Mika Maliranta. Globalization, creative destruction, and labour share change: Evidence on the determinants and mechanisms from longitudinal plant-level data. *Oxford Economic Papers*, 64(2):259–280, April 2012. (Cited on page 7.)

Christoph E. Boehm, Aaron Flaaen, and Nitya Pandalai-Nayar. Multinationals, offshoring, and the decline of U.S. manufacturing. *Working Paper*, 2015. (Cited on page 7.)

Benjamin Bridgman. Is labor's loss capital's gain? gross versus net labor shares. *Working Paper*, 2014. (Cited on page 8.)

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 8.)

Russell W. Cooper, John C. Haltiwanger, and Jonathan L. Willis. Declining dynamism at the establishment level. *SED Meeting Paper*, 2017. (Cited on page 40.)

Steven J. Davis, John C. Haltiwanger, and Scott Schuh. *Job Creation and Destruction*. MIT Press, Cambridge, MA, 1996. (Cited on page 45.)

Jan De Loecker, Jan Eeckhout, and Gabriel Unger. The rise of market power and the macroeconomic implications. *Working Paper*, November 2018. (Cited on pages 8 and 11.)

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 39.)

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 page 40.)

Robert C. Dent, Benjamin W. Pugsley, and Harrison Wheeler. Longitudinal linking of enterprises in the LBD and SSL. *CES Technical Notes CES-TN-2018-02*, 2018. (Cited on page 45.)

Matthew Dey, Susan N. Houseman, and Anne E. Polivka. Manufacturers' outsourcing to staffing services. *ILR Review*, 65(3):533–559, July 2012. (Cited on page 46.)

Maya Eden and Paul Gaggl. On the welfare implications of automation. *Review of Economic Dynamics*, 29:15–43, July 2018. (Cited on pages 7 and 13.)

Chris Edmond, Virgiliu Midrigan, and Daniel Y. Xu. How costly are markups? *NBER Working Paper No. 24800*, 2018. (Cited on pages 8, 11, and 33.)

Andrea Eisfeldt, Antonio Falato, and Mindy X. Zhang. Human capitalists. *Working Paper*, 2018. (Cited on page 47.)

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 7, 8, 12, and 13.)

Rudy Fichtenbaum. Do unions affect labor's share of income: Evidence using panel data. *American Journal of Economics and Sociology*, 70(3):784–810, July 2011. (Cited on page 12.)

Teresa Fort and Shawn Klimek. The effects of industry classification changes on US employment composition. *Census Discussion Paper CES 18-28*, June 2018. (Cited on page 46.)

Teresa Fort, Justin R. Pierce, and Peter K. Schott. New perspectives on the decline of US manufacturing employment. *Journal of Economic Perspectives*, 32(2):47–72, Spring 2018. (Cited on page 18.)

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 27, 31, 32, and 48.)

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 7.)

Andrew Glover and Jacob Short. Demographic origins of the decline in labor's share. *Working Paper*, 2018. (Cited on pages 7 and 12.)

Emilien Gouin-Bonenfant. Productivity dispersion, between-firm competition and the labor share. *Working Paper*, 2018. (Cited on pages 12 and 17.)

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 pages 7 and 11.)

Gustavo Grullon, Yelena Larkin, and Roni Michaely. Are US industries becoming more concentrated? *Working Paper*, 2016. (Cited on pages 8 and 11.)

Germán Gutiérrez and Thomas Philippon. Declining competition and investment in the U.S. *NBER Working Paper No. 23583*, 2017. (Cited on page 11.)

Barney Hartman-Glaser, Hanno N. Lustig, and Mindy X. Zhang. Capital share dynamics when firms insure managers. *Journal of Finance*, 74(4):1707–1751, August 2019. (Cited on page 7.)

Brad Hershbein, Claudia Macaluso, and Chen Yeh. Labor market concentration and the demand for skills. *Working Paper*, 2018. (Cited on page 12.)

Susan N. Houseman. Understanding the decline of U.S. manufacturing employment. *Upjohn Institute Working Paper 18-287*, January 2018. (Cited on page 46.)

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 27.)

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 40.)

Gregor Jarosch, Jan Sebastian Nimczik, and Isaac Sorkin. Granular search, market structure, and wages. *NBER WP No. 26239*, September 2019. (Cited on page 12.)

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 5.)

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 7 and 11.)

Loukas Karabarbounis and Brent Neiman. Capital depreciation and labor shares around the world: Measurement and implications. *Working Paper*, 2014b. (Cited on page 8.)

Barış Kaymak and Immo Schott. Corporate tax cuts and the decline of the labor share. *Working Paper*, 2018. (Cited on pages 7 and 11.)

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 8 and 45.)

Miles S. Kimball. The quantitative analytics of the basic neomonetarist model. *Journal of Money, Credit and Banking*, 27:1241–1277, 1995. (Cited on page 11.)

Dongya Koh, Raül Santaeulàlia-Llopis, and Yu Zheng. Labor share decline and intellectual property products capital. *Working Paper*, 2019. (Cited on pages 7 and 13.)

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 7.)

Aslı Leblebicioğlu and Ariel Weinberger. Credit and the labor share: Evidence from U.S. states. *Working Paper*, 2018. (Cited on page 11.)

Marc J. Melitz and Gianmarco I. P. Ottaviano. Market size, trade, and productivity. *Review of Economic Studies*, 75(1):295–316, January 2008. (Cited on page 11.)

Brent Neiman and Joseph Vavra. The rise of niche consumption. *NBER Working Paper No. 26134*, August 2019. (Cited on pages 8 and 11.)

Ezra Oberfield and Devesh Raval. Micro data and macro technology. *NBER Working Paper No. 20452*, 2014. (Cited on page 7.)

G. Steven Olley and Ariel Pakes. The dynamics of productivity in the telecommunications equipment industry. *Econometrica*, 64(6):1263–1297, November 1996. (Cited on page 9.)

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 13.)

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 49.)

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 27, 32, and 48.)

Appendix

A Data and measurement

A.1 Constructing the Full Census Sample

Say sth about establishments vs. firms (Ref 1) and outsourcing.

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 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.3) and value added (described in Section A.4).

A.2 Identifying establishments, firms and industries

All establishments carry an identifier, `LBDNUM`, which stays with the establishment from its birth to its death. That variable is available as a consistent identifier throughout all years. In addition to that, every establishment carries a firm identifier, `FIRMID`, which owns the establishment.[21] Unlike the `LBDNUM`, the `FIRMID` may change over time, especially when a firm transitions from a single-unit to a multi-unit firm and vice versa (see Dent, Pugsley, and Wheeler (2018)). We account for that possibility when we study firm-level dynamics.

We identify an establishment's industry from its SIC code (until 1996) and then its NAICS code. We map SIC codes into NAICS codes as in Kehrig (2011) and consider only establishments active in manufacturing industries (NAICS code 311111 through 339999). This entails first correcting for erroneous industry classifications 1972 to 1986 according to the list on p. 222 in Davis, Haltiwanger, and Schuh (1996). Then, SIC-72 codes were mapped into SIC-87 codes. In case of non-unique mappings, we settle on the SIC-87 industry which captures most of the employment of the SIC-72 industry. SIC-87 industries are mapped into NAICS industries using concordance files provided by the Census Bureau. Whenever this mapping didn't produce a unique industry code, we used an establishment's NAICS code as sampled rather than the one implied by the SIC-NAICS concordance. Discrepancies may occur between the two when establishments predominantly tasked with corporate activities were initially labeled as a manufacturing and later as a services establishment. Picking the sampled NAICS code (and dropping non-manufacturing establishments) makes

[21] In case of joint ownership, this appears to be the firm owning the majority stake in the establishment.

our procedure similar in spirit to that in Fort and Klimek (2018), although these authors deal with that matter (and other problems) more comprehensively than we do. Note that some of these industry changes from manufacturing to services may actually be legitimate because establishments that used to perform predominantly production activities may transition into a support-activity establishment. As a robustness check, we also use their industry codes to verify our main findings. Both our industry way to consistently identify industry codes as well as the Fort-Klimek codes yield similar results for labor share dynamics even though Fort and Klimek (2018) document strong differences for employment dynamics.

A.3 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. We denote their sum by the variable `LC`.

To properly measure an establishment's labor share, we have to sum up the compensation of all employees that help generate the establishment's value added. In principle, Census attempts to do that, but it is not certain that all temporary help services, or leased employment, was captured in the earlier Census years before 2002. In the past decades, leasing workers rather than employing them full-time has become increasingly popular in U.S. manufacturing (see Dey, Houseman, and Polivka (2012); Houseman (2018)). To accommodate that important new development, Census started to sample permanent and leased employment separately starting in 2002. Before then, no specific instructions were given to establishment whether or not to include leased employment in their compensation variables. After studying the before/after patterns of employment and labor costs, Census believes that the majority of establishments interpreted the question to include all types of workers and their compensation (see **?**). If true, this would be ideal in the context of our work, and our labor share measures would capture all labor cost before and after 2002. If not true, we would miss a portion of the labor compensation in 1997 and before, so we would underestimate labor shares in that early time period. Given Census's before/after analysis this missing labor compensation is likely small. But even if it was significant, the manufacturing labor share decline would be even stronger than what we measure, so we can view our empirical results as a lower bound on the manufacturing labor share decline.

Put cite!

To summarize, we measure labor compensation as follows:

- Before 2002: `SW + LC`, which supposedly comprises salaries, wages and benefits for both permanent and leased employees;

- 2002: `SW_NL + BENEFIT_NL + SW_L + BENEFIT_L`. The first two terms consist of salaries, wages and benefits for permanent (non-leased) employees; the latter two consist of their analogues for temporary (leased) employment.

- 2007 and later: SW + BENEFIT + CTEMP. The first two terms consist of salaries, wages and benefits for permanent employees only (note how SW now captures a subset of what it used to capture before 2002); the last term combines all compensation for temporary (leased) employment. (In 2007 and later, Census does not parse out the cost of leased employment into salaries, wages and benefits as it did in 2002).

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.[22] 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 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 manufacturing labor share decline.

A.4 Measuring value added

We measure value added in the Census data 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), material inputs (sum of items CP, CW and MIB less MIE) and energy expenditures (sum of items CF and EE). This procedure refines the definition of value added vis-a-vis the previous literature using the standard Census definition in two ways:

1. Materials use is corrected for inventory adjustment.
2. Industry-year-specific measures of purchased services are added to intermediate input use.

Both steps bring our measure closer to the value actually added by the establishment as a manufacturer.

ad 1. Constructing value added requires subtraction of all materials inputs regardless if they were purchased in the same period or came out of the materials inventory. Failure to do so would make both value added and revenue total factor productivity (TFPR) too volatile over time and too dispersed across establishments because they would both include a portion of unmeasured fluctuations in intermediate inputs. Since value added plays an important role in the dynamics and aggregation of labor shares, this matters. Getting only biased measures of TFPR would cause problems in Sections 4.5 through 4.6.

Also, note that establishments where inventory changes play a major role essentially act as a warehouse and may eventually be classified as such (NAICS code 4931XX), thus leading to discontinuous jumps in aggregates that inaccurately reflect the true economic activity.

ad 2. The Census of Manufactures samples intermediate energy and material inputs as well as contract work, but information about an establishment's purchased services is absent. This makes value added too large and the labor share too low. As a consequence, the raw aggregate manufacturing labor share in Census data is about 14 ppt lower than its BLS counterpart

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

in 1967. Importantly, this discrepancy gets worse over time because outsourcing of non-manufacturing activity and purchased services grew substantially over the past decades. As a result, the raw aggregate manufacturing labor share in Census data is 20 ppts lower than its BLS counterpart in 2012, thus overstating the decline in the manufacturing labor share.

We therefore subtract the industry-year-specific ratio of purchased services to sales from establishment sales.[23] This avoids outsourcing contaminating our measure of the labor share and its time series behavior, but it does not impact within-industry reallocation dynamics because that correction is identical for all establishments in a given industry and year. Correcting for the increasing prevalence of purchased services in this way reduces the overall difference between Census and BLS manufacturing labor shares to 5 ppt, a difference that stays constant over time (see Figure **??**).

A.5 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,
  - `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;

[23] Autor, Dorn, Katz, Patterson, and Reenen (2017b) pursue a similar strategy.

– `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.6 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 LL and Non-LL establishments exist.

In order to produce Figure A.1, we run a non-parametric regression analogous to Equation (17) 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 LL establishments versus Non-LL establishments. The approach is analogous to (13) and (14): 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 LL 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 LL establishments jump relative to the Non-LL establishments by 21% during the five years preceding the year in which they become LL. 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 LL vs Non-LL establishments. Over the

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

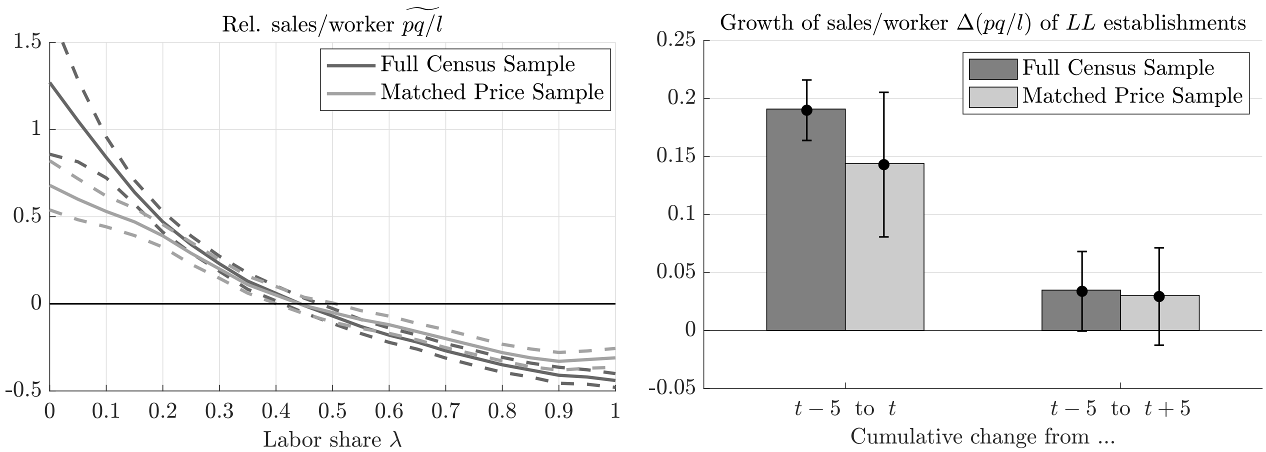

Note: The left panel in the figure depicts the cross-sectional differences in relative sales per worker $\widetilde{pq/l}$ against the labor share in the Full Census Sample (dark grey line) and the Matched Price Sample (light grey line). Dashed lines denote 95% error bands.
The right panel displays the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

entire time span, the estimates for the Full Census Panel show a significantly different sales per worker trajectory for LL establishments than for Non-LL 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 LL status, sales per worker grow by 12.5% more for LL 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 The roles of industry and regional composition

To test for industry and/or geographical composition effects, we decompose the manufacturing labor share decline into within- and between-groups components using Equation (21):

$$\Delta\lambda_t = \underbrace{\sum_j \Delta\lambda_{jt}\omega_{jt-5}}_{\text{Within adjustment}} + \underbrace{\sum_j \lambda_{jt-5} d\omega_{jt}}_{\text{Between reallocation}} + \underbrace{\sum_j \Delta\lambda_{jt} d\omega_{jt}}_{\text{Residual}} \tag{21}$$

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

Panel A. in Table B.1 displays the results from an 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.

Turning our attention to the regional dimension, Panel B. in Table B.1 shows that 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.[24]

Table B.1: 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)		
Manufacturing 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

Note: Results from the shift-share decompositions as defined in (21) 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.

B.2 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.

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

Production worker wages include the wage bill of all employees engaged in the core manufacturing activities such as fabricating, processing, assembling, inspecting, receiving, packing, warehousing, maintenance, repair, janitorial and guard services and record keeping. Salaries of non-production workers refer instead to the compensation of all employees above line-supervisor level; it comprises executive, purchasing, professional and technical sales, logistics, advertising, credit, clerical and routine office functions. Finally, the ancillary labor costs comprise legally-required labor costs (such as social security tax, unemployment tax, workmen's compensation insurance and state disability insurance pension plans) as well as voluntary labor costs (such as health benefits, life insurance premiums, supplemental unemployment compensation and deferred profit sharing plans).

We investigate 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 B.2 and Table B.2:

$$\lambda_t = \underbrace{\frac{w_t^{pw} L_t^{pw}}{Y_t}}_{\text{Wage bill}} + \underbrace{\frac{w_t^{npw} L_t^{npw}}{Y_t}}_{\text{Salaries}} + \underbrace{\frac{w_t^{ben} L_t^{ben}}{Y_t}}_{\text{Ancill. labor costs}} . \tag{22}$$

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

Figure B.2: Dynamics of labor share components

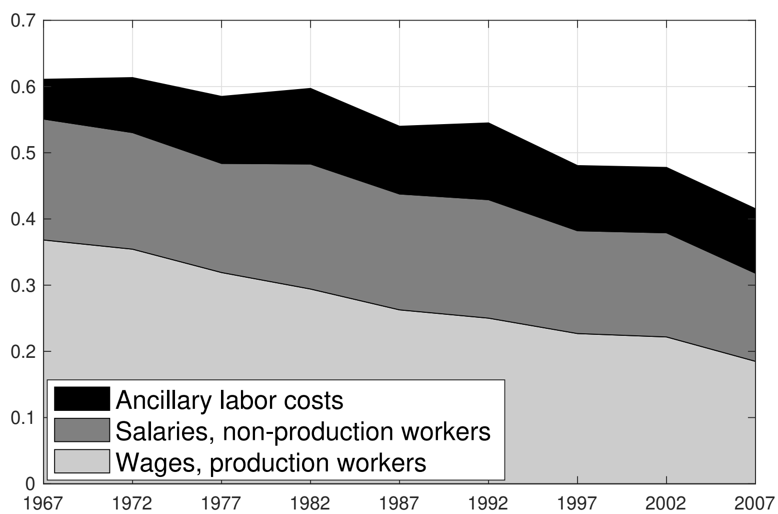

Note: This figure displays three portions of manufacturing labor compensation: wages of production workers, salaries of non-production workers and total ancillary labor costs such as unemployment insurance and health benefits. The secular decline of the production worker wage bill was first compensated 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 manufacturing 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 manufacturing labor share would have stayed more or less constant (-0.3 ppt per decade).

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

Component	1967-2007	1967-1982	1982-2007
	(percentage point changes)		
Manufacturing 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 (21) applied to the three types of labor compensation listed in Equation (22). 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.

The compensation for non-production labor, in contrast, is steady at first and then starts to 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 manufacturing 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 manufacturing 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 manufacturing 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 manufacturing labor share decline would have been stronger at 5.8 ppt per decade instead of the observed 4.9 ppt decline.

B.3 Transitory versus permanent LL establishments

In Section 4.4 we showed that LL 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 LL establishments do have a permanently low labor share and are among the LL 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" LL establishments. Since the former tend to be larger and thus more relevant for aggregates, we want to ensure that the "temporary LL establishments," those characterized by the V-shaped pattern of Figure 7, play a significant role for the manufacturing labor share decline.

To that end, we partition the set of LL establishments in period t into those that are an LL establishment from $t-4$ to $t+5$, denoted "permanent LL," and the rest, denoted "temporary LL." When we drop both temporary and permanent LL establishments from the sample, the manufacturing labor share has a much higher level and stagnates. This shows that LL establishments are essential to understanding the manufacturing labor share decline; see the light grey line in Figure B.3. When we instead only drop the permanent LL 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 LL establishments play an important role for the manufacturing labor share level and its decline.

Fig B.3 is coming in Revision.xls > LS_weighted

Figure B.3: The role of temporary and permanent LL establishments

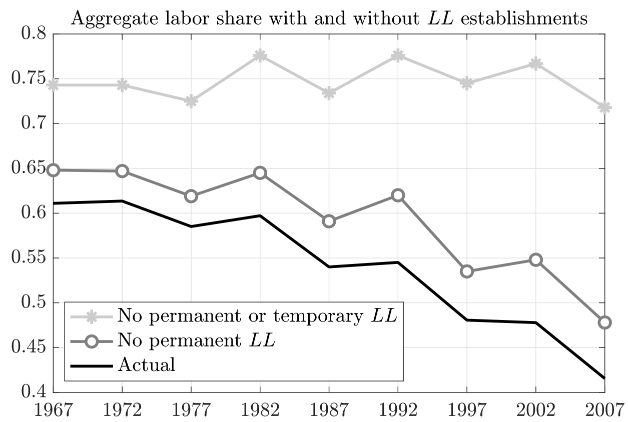

Put the LL vs. LL2 result here. Pareto distribution and Markov matrix

B.4 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 LL establishments likely sort into the same firms. Labor shares of firms that operate mostly LL establishments should then exhibit the same V-shaped pattern that we observe for the LL establishments in Figure 7. If, on the other hand, LL 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. In principle, `FIRMID`s stay with the same firms although firms that transition from single-unit to multi-unit firms receive a new type of `FIRMID` even though the firm is not new. For sake of simplicity, we perform a simple robustness check in which we drop those observations from the data and verify that the main results about firm displayed here do not change.

Analogously to LL establishments, we define "LL 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 (13) and (14) for these LL firms and show them in Figure B.6. 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 4.6. Second, LL establishments tend to assort into the same firms.

Nic, what else should we disclose about firms?

Fig B.6 is coming in New.xls > VShape_Firms

Figure B.4: LL establishments versus LL firms

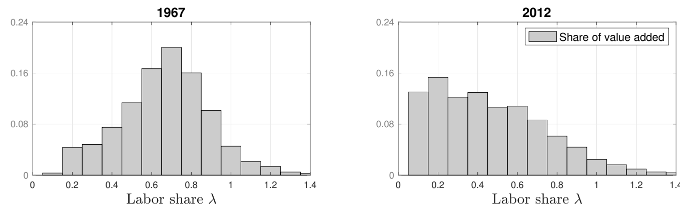

Figure B.5: LL establishments versus LL firms

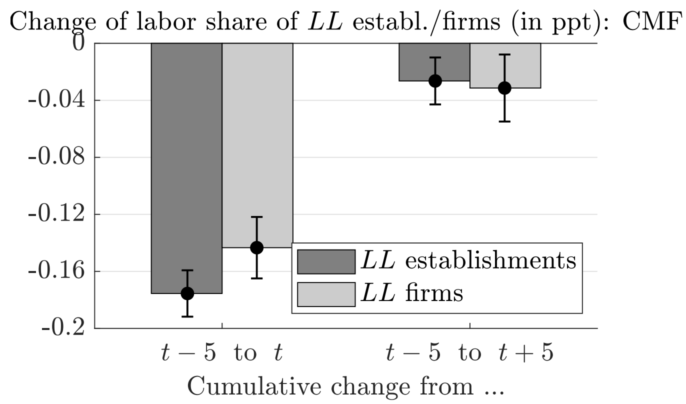

Figure B.6: LL establishments versus LL firms

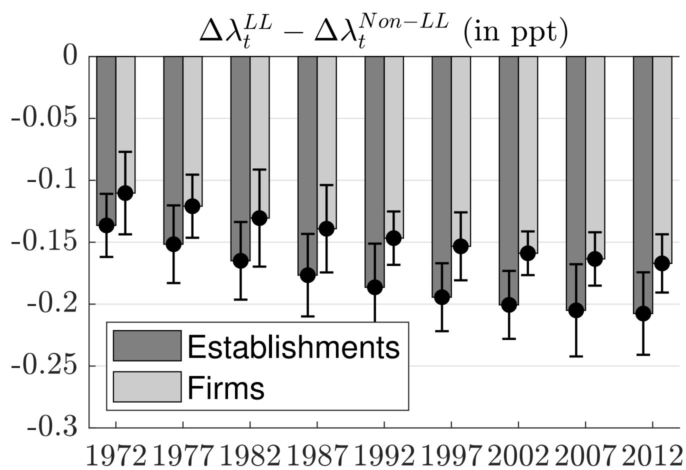

Contribution of rel. growth in L, W and Y to $\Delta\lambda^{LL}$ (in ppt)

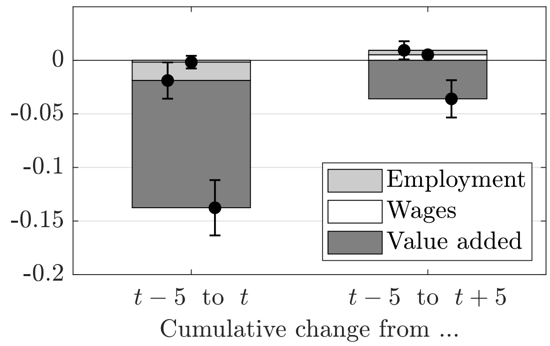

Contribution of labor, value added and wage growth: 2000s

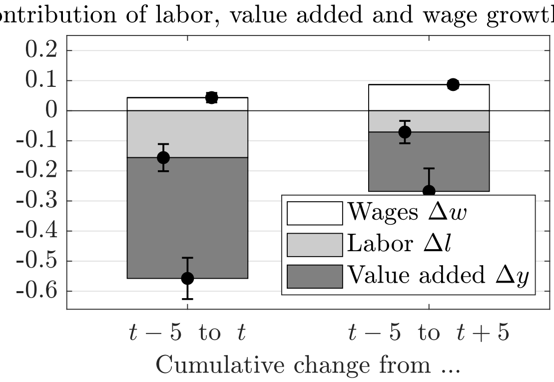

Contribution of labor, value added and wage growth: 1970s

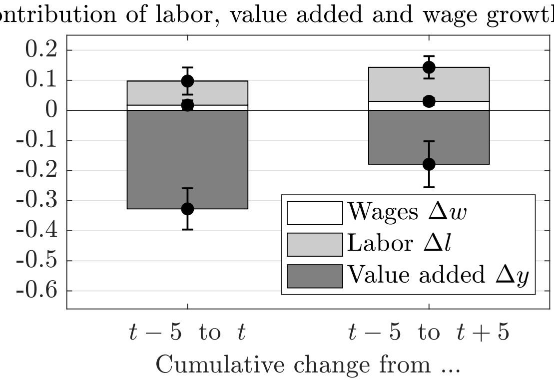

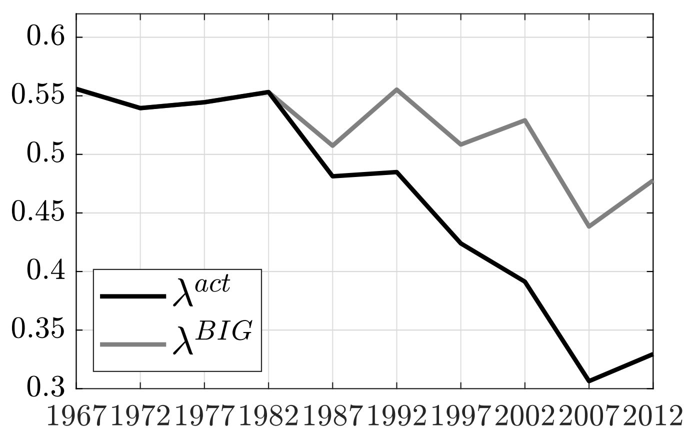

0.6
0.55
0.5
0.45
0.4
0.35
0.3
1967
1972
1977
1982
1987
1992
1997
2002
2007
2012
λ^{act}
λ^{BIG}

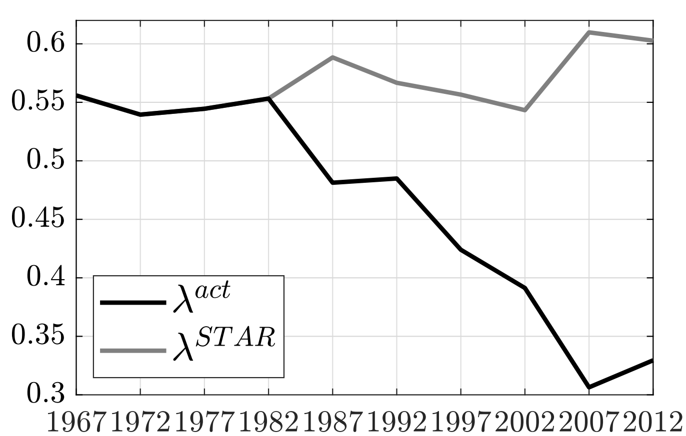

0.6
0.55
0.5
0.45
0.4
0.35
0.3
1967
1972
1977
1982
1987
1992
1997
2002
2007
2012
λ^{act}
λ^{STAR}