

The Micro-Level Anatomy of the Labor Share Decline

by

Matthias Kehrig
Duke University

Nicolas Vincent
HEC Montreal

CES 20-12

March, 2020

The research program of the Center for Economic Studies (CES) produces a wide range of economic analyses to improve the statistical programs of the U.S. Census Bureau. Many of these analyses take the form of CES research papers. The papers have not undergone the review accorded Census Bureau publications and no endorsement should be inferred. Any opinions and conclusions expressed herein are those of the author(s) 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. Republication in whole or part must be cleared with the authors.

To obtain information about the series, see www.census.gov/ces or contact Christopher Goetz, Editor, Discussion Papers, U.S. Census Bureau, Center for Economic Studies 5K038E, 4600 Silver Hill Road, Washington, DC 20233, CES.Working.Papers@census.gov. To subscribe to the series, please click [here](#).

Abstract

The labor share in U.S. manufacturing declined from 62 percentage points (ppts) in 1967 to 41 ppts in 2012. The labor share of the typical U.S. manufacturing establishment, in contrast, rose by over 3 ppts during the same period. Using micro-level data, we document five salient facts: (1) since the 1980s, there has been a dramatic reallocation of value added toward the lower end of the labor share distribution; (2) this aggregate reallocation is not due to entry/exit, to “superstars” growing faster or to large establishments lowering their labor shares, but is instead due to units whose labor share fell as they grew in size; (3) low labor share (LL) establishments benefit from high revenue labor productivity, not low wages; (4) they also enjoy a product price premium relative to their peers, pointing to a significant role for demand-side forces; and (5) they have only temporarily lower labor shares that rebound after five to eight years. This transient pattern has become more pronounced over time, and the dynamics of value added and employment are increasingly disconnected.

Keyword: Labor Share, Productivity, Firm Size Distribution, Relative Prices

JEL Classification: E2, L1, L2, L6, O4

* We would like to thank Gadi Barlevy, Nick Bloom, 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, Jeff Campbell, Chris Gust, Ezra Oberfield and Kirill Shakhnov as well as conference and seminar participants at many places 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 Montreal (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.
Kehrig: Duke University, Department of Economics, NBER and CEPR. Email: matthias.kehrig@gmail.com.
Vincent: HEC Montreal, Department of Applied Economics. Email: nicolas.vincent@hec.ca.

1 Introduction

Several recent studies have documented a decline of the aggregate labor share, the portion of gross domestic product paid out in compensation for labor. This finding is important for several reasons. For one, it contradicts one of the stylized facts of [Kaldor \(1961\)](#) that have become foundational for economic growth theories. It is further at odds with a key building block of standard macroeconomic models, the Cobb-Douglas production function. Last, the finding 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.

Numerous explanations have been suggested to explain this aggregate decline, most of which are rooted in firm-level behavior, yet little is known about the dynamics at the micro level and which structural drivers and shocks are consistent with the empirical picture. This paper fills this gap. It studies the micro-level anatomy of labor shares and factor reallocation, which may serve as a basis for future modeling work. We first present a conceptual framework that encompasses three of the leading theories that have been proposed in the literature: demand/pricing factors, total factor productivity (TFP)/efficiency channels and market power in labor markets. We then use confidential data from the U.S. Census of Manufactures (CMF) to study the establishment- and firm-level anatomy of labor shares, with the aim of identifying those theories consistent with the empirical evidence. Based on our empirical work, we argue that only demand factors are capable of explaining both the micro anatomy of labor shares and factor allocations as well as the manufacturing labor share decline.

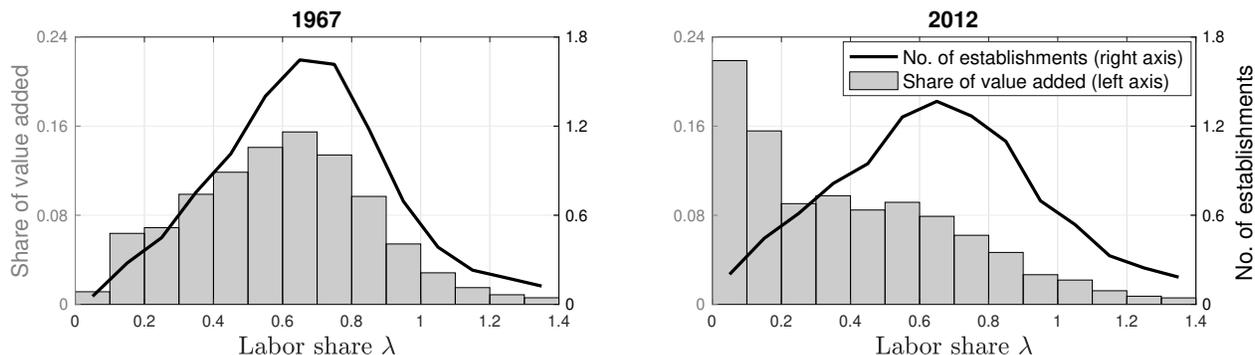
We document a number of salient facts. First, we confirm that the labor share in the manufacturing sector declines by almost 5 percentage points (ppts) per decade between 1967 and 2012. This decline, however, hides contrasting dynamics at the micro level: alongside the sectoral decline, the median establishment saw an *increase* in its labor share, by about 0.7 ppts per decade. In fact, this upward trend is present for the vast majority of manufacturing establishments. We find that the decline of the manufacturing labor share is entirely driven by a strong reallocation of value added toward the low end of the labor share distribution (see [Figure 1](#)). In other words, the covariance between establishment labor shares and establishment value-added weights turned sharply negative over time. In contrast, the reallocation of labor was much less pronounced over the same period.

Second, we show that this reallocation was not driven by entry and exit, by large establishments lowering their labor shares over time, nor by superstars with initially low labor shares increasing their market share. 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, focusing on the components of labor shares, we show that cross-sectional differences are almost entirely driven by value added per worker, not wages. We then focus on establishments in the bottom quintile of the labor share distribution in a specific year and industry, which we define as low labor share (*LL*) establishments. We find that the labor share dynamics of *LL* establishments are shaped by value added, with very little accounted for by employment or wages.

Fourth, we find that low labor share establishments tend to have, on average, significantly higher

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 gray bars (left axis). The labor share on the x-axis is expressed as a decimal.

prices than their peers. Moreover, price dynamics are first order in understanding the dynamics of sales per worker for *LL* establishments. We reach this conclusion by using a subsample of the census database that provides information about the value of sales and physical quantity for individual products. Doing 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.

Fifth, we find that labor share fluctuations at the establishment level are surprisingly transient: the probability that a typical *LL* establishment today loses its *LL* status five years from now is close to 60%. This number would be close to 0% if *LL* establishments had permanently low labor shares. Even more surprisingly, we document that the labor share dynamics of *LL* establishments follow a 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. In that sense, *LL* establishments are more like “shooting stars” than “superstars.”

We complete our empirical analysis by highlighting the evolution of two central features of the micro-level anatomy of the labor share. We first show that the V-shaped labor share pattern described earlier has become more pronounced over time. Moreover, we document that employment has become increasingly disconnected from variations in value added.

These findings, which hold if we instead use firm-level data, have important implications for our understanding of the drivers of labor share dynamics. In the context of our conceptual framework, they make a strong case for a significant role played by demand factors: technology shocks would counterfactually predict that prices drop with labor shares, while monopsony power in labor markets would imply a counterfactual relationship between wages and labor shares. These demand factors could take many forms. For example, they may be driven by strong, but ephemeral, brand appeal or by fads for superstar products. Alternatively, firms could spur higher demand for their products through advertising campaigns, which may have become farther-reaching and more effective with

the advent of online marketing and the overall growth of market scope. Such developments would be compatible with our evidence pointing to an increase in the volatility of the demand factors that underlie the micro-level dynamics of the labor share.

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

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 \(2014a\)](#) puts forward the notion that technical change embodied in new equipment capital has displaced labor and has 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, Santaaulàlia-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 \(2018\)](#) 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 United States. In related work, [Boehm, Flaaen, and Pandalai-Nayar \(2015\)](#) present establishment-level evidence that outsourcing did cut U.S. manufacturing employment while raising the profits per worker of surviving production units. [Glover and Short \(2018\)](#) find the workforce’s age composition has shifted toward workers who 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 Organisation for Economic Co-operation and Development (OECD) countries.

[Furman and Orszag \(2015\)](#) notes that the distribution of capital returns—inversely related to the labor share—had shifted up and became more skewed toward 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 the 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, in turn raising 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 that requires more ex ante insurance.

Last, an emerging strand of the labor share literature emphasizes the role of rising concentration and markups. [Autor, Dorn, Katz, Patterson, and Van Reenen \(2017, forthcoming\)](#), for example, present some industry- and firm-level evidence on 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 engaging in 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. As [Edmond, Midrigan, and Xu \(2018\)](#) show, the rise in markups largely disappears if firm-level markups are aggregated with the proper weights. They nevertheless find large welfare implications of high markups and high-markup dispersion and 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. [Baqee and Farhi \(2020\)](#) 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, *LL* 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 *LL* 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 the large sales 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 worldwide 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 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:

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

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

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, $\bar{\lambda}_{it}$, 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 $\bar{\lambda}_{it}$.

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 wage bill while keeping 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. Last, 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 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

$$\lambda_{it} = \frac{W_{it}L_{it}}{P_{it}Y_{it}} = \frac{W_{it}}{ARPL_{it}} = \frac{W_{it}}{P_{it}APL_{it}}, \quad (3)$$

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} :

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

$$\omega_{it} = \frac{P_{it}Y_{it}}{\sum_i P_{it}Y_{it}} = \frac{\mu_{it}MC_{it}Y_{it}}{\sum_i P_{it}Y_{it}}, \quad (5)$$

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

where μ_{it} corresponds to the price markup, 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 i 's products or brand image more 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 as in [Kimball \(1995\)](#) or [Melitz and Ottaviano \(2008\)](#), 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, e.g., [Grullon, Larkin, and Michaely \(2016\)](#); [Barkai \(2017\)](#); [Baqee and Farhi \(2020\)](#); [Gutiérrez and Philippon \(2017\)](#); [Edmond, Midrigan, and Xu \(2018\)](#); [De Loecker, Eeckhout, and Unger \(2018\)](#); [Neiman and Vavra \(2019\)](#).

Supply-side 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 \(2018\)](#) and [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, producers 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 Van Reenen \(2017, forthcoming\)](#). 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 \(forthcoming\)](#); [Karabarbounis and Neiman \(2014a\)](#). Notice that these dynamics are similar to those under the scenario of

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

a demand shock with non-isoelastic demand schedules, except for one important difference: prices will decline after supply-side TFP shocks, while they will increase under demand shocks.

Monopsony power Last, let us turn to the role of market 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 markdown, 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} .

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

$$\omega_{it} = \frac{P_{it}Y_{it}}{\sum_i P_{it}Y_{it}}. \quad (7)$$

Note that by itself, a lower ν_{it} decreases the unit’s labor share but does not increase its value-added weight unless it is profitable to expand scale and/or adjust its price relative to its peers. A stronger wage markdown 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 Sorokin \(2019\)](#)), labor market deregulation, such as de-unionization ([Fichtenbaum \(2011\)](#)), “right-to-work” laws ([Blanchard and Giavazzi \(2003\)](#)), demographic factors ([Glover and Short \(2018\)](#)) or search-and-matching frictions ([Gouin-Bonenfant \(2018\)](#)). While the empirical evidence on the link between business concentration trends and labor shares is ambiguous (see [Berger, Herkenhoff, and Mongey \(2019\)](#); [Hershbein, Macaluso, and Yeh \(2018\)](#)), the use of micro-level data allows 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 of 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 manufacturing 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, making 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 CMF 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 allows us 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. Sixth, 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, Santaaulàlia-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 CMF 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 the Census Bureau. 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

$$\lambda_t = \frac{W_t L_t}{P_t Y_t}, \quad (8)$$

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

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

provides us the advantage of avoiding measurement issues related to inflation.

We define the following items as labor costs: salaries and wages for permanent and leased workers, involuntary labor costs such as unemployment insurance or social security contributions netted out from wages and voluntary labor costs 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. We also drop all observations in the bottom and top percentiles to avoid outliers driving our results. In doing so, 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 is that any theory of manufacturing labor share dynamics should be compatible with these salient facts. Though our analysis is at the establishment level, all subsequent results also hold if we aggregate to the firm level. We present those firm-level results in Appendix C.

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 Equations (1) and (2) and that we reproduce in the first line below:

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

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

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

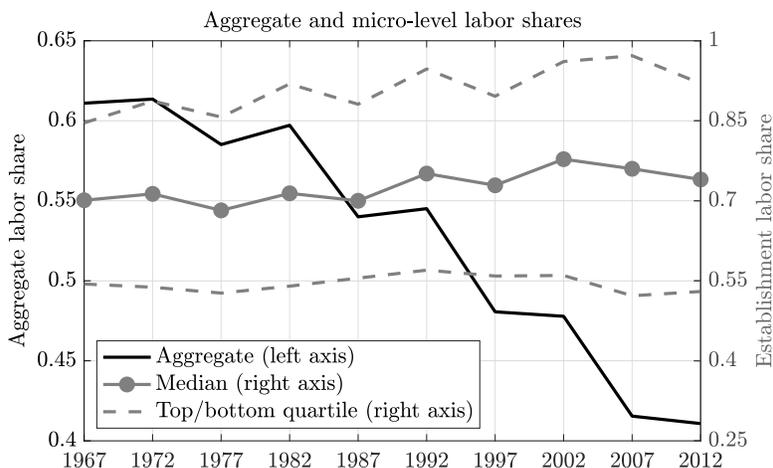
From this decomposition, we can readily identify two broad ways a decline in the manufacturing labor share may come about. First, it could follow from a general decline of the unit-level labor shares λ_{it} , which would be reflected in a lower (unweighted average) $\bar{\lambda}_{it}$. This may, for 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.

⁴The Census Bureau does not collect information on nonmonetary 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. Such an increase in unmeasured compensation could potentially mitigate the manufacturing labor share decline.

4.2.1 The Labor Share of the Median Establishment *Increases*

We now aim to disentangle these various scenarios with the help of micro-level data. As a first exercise, 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: 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 (gray lines, right axis): the solid gray line with circles reflects the “median,” per U.S. Census disclosure rules, defined as the average of the sample of observations between the 49th and 51st percentile; the dashed gray lines reflect the first and third “quartile,” defined analogously to the “median.” While the manufacturing labor share declines strongly, the “median” and top “quartile” labor share increase over time.

Figure 2 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 ppts per decade, on average, the median labor share *increases* by 0.7 ppts 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.⁵ This is what we turn our attention to next.

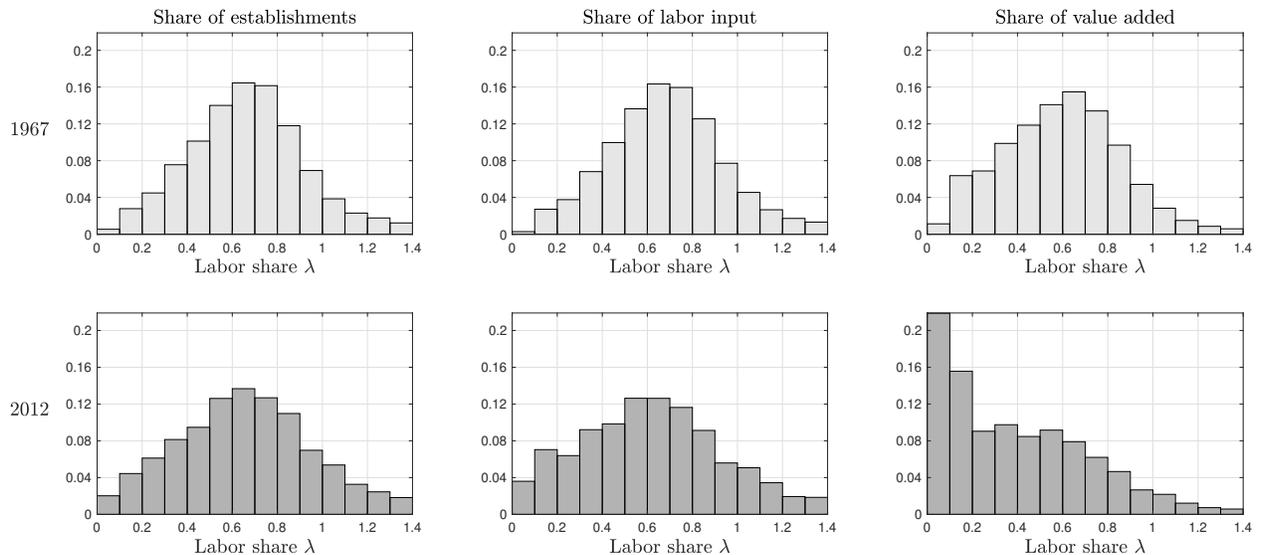
4.2.2 Reallocation: Dramatic for Value Added, Anemic for Labor

The diverging trends between macro- 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 toward the left 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 ppts

⁵In Appendix A.3, we show that the decline in the manufacturing labor share is present for both production and nonproduction workers.

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 compute these shares for each three-digit NAICS industry and then aggregate them up in each bin using the industry’s value-added weight at the annual level. The subsequent analysis therefore refers to reallocation of value added *within* a typical industry.⁶

Figure 3: Distribution of Establishments, Labor Input and Value-Added Conditional on Labor Share



Note: The bars in the two panels in the first column reflect the raw cross-establishment distribution of labor shares in 1967 (light gray on top) and 2012 (dark gray on bottom). The panels in the middle column display the allocation of labor, and those in the right column display that of value added. To account for industry-specific differences in the distributions, we first calculated the distributions within each three-digit NAICS industry. Then these industry-specific 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 establishments, hours worked and value added in that year.

The light gray bars in the top row of panels in Figure 3 display the distributions of the number of establishments (left), labor input (middle) and value added (right) against the labor share in 1967; the dark gray bars in the bottom row show the analog distributions in 2012. The unweighted distribution of establishments against the labor share (panels in left column) did not change location, which confirms the evidence of the largely stagnant quantiles of the unweighted distribution in Figure 2. There appears to be a slight fattening of the tails, also visible in Figure 2. The distribution of labor input in the middle column panels is measured as employment here, but the distributions of production hours worked look nearly identical. They suggest only a minute reallocation of labor input to low labor share establishments. But the panels in the right column paint a stark picture of drastic reallocation of output: most of value added in 1967 is produced by establishments in the

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

middle of the labor share distribution (between 50 and 80 ppts). The value-added weighted median labor share is 62 ppts. Over the following decades, however, the economic activity shifts gradually and persistently toward the low labor share spectrum: by 2012, half of manufacturing value added is accounted for by establishments with a labor share less than 32 ppts. The disconnect between value added and labor reallocation is a key feature of the labor share decline. Similar evidence at the firm level has been found for other sectors in the U.S. by [Autor, Dorn, Katz, Patterson, and Van Reenen \(forthcoming\)](#), for Canada by [Gouin-Bonenfant \(2018\)](#) and for China by [Berkowitz, May, and Nishioka \(2017\)](#).

Referring back to our discussion surrounding Equation (2), Finding 1 makes it clear that common trends (e.g., due to a generalized increase in markups or monopsony power) are unlikely to be behind the decline in the manufacturing labor share. Such a general development would have manifested itself through a leftward shift of the unweighted distribution by 20 ppts, on average, yet it has remained centered around $\lambda = 0.65$ (see left column in Figure 3). Hence, the manufacturing labor share decline must be driven by a strong decline in the covariance between establishment-level labor shares and market shares: since the 1980s, low labor share establishments—though small in number—have also happened to be much larger producers than their high labor share peers, as is visible in the right column of Figure 3. In contrast, the middle column indicates that the distribution of the labor input did not follow the same dramatic pattern: this concentration of value added did not come with a similar shift in the distribution of employment. 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 reveal how the reallocation of value added came about. Consider three very different theories that would generate a fall in the manufacturing labor share through a decline in $Cov(\lambda_i, \omega_i)$. First, entry/exit could be behind the sectoral 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 of value added, $\Delta\omega_{it}$, can have an aggregate impact if they are correlated with the initial 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 (e.g., 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 shares and market shares that we document in Finding 1. In this section, we put them to the test with the help of our detailed data.

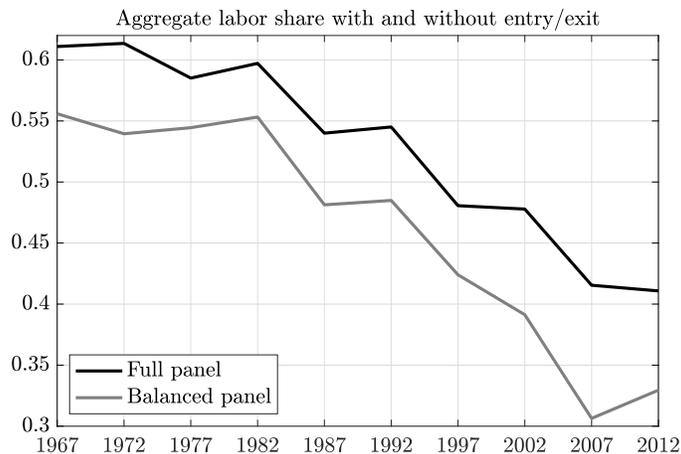
4.3.1 The Negligible Role of Entry and Exit

First, we test whether the decline in the manufacturing labor share is driven by the 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 of establishments that are permanently active from 1967 to 2012.⁷ If entry and exit were the primary drivers of the labor share decline, the labor share in this strongly balanced sample should be stagnant.

Figure 4: Manufacturing Labor Share in Full and 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 gray line). It shows that entry and exit matter for the *level* but not for the *decline* of the labor share.

Figure 4 reveals that the manufacturing labor share in the strongly balanced sample is, on average, about 6 ppts lower than that of the full sample. Yet, the labor share trends in the full and balanced samples are virtually indistinguishable. If anything, the decline is even more pronounced if we ignore entry and exit: the manufacturing labor share in both samples stagnates until 1982 and then falls by 6.2 ppts per decade from 1982 to 2012 in the full sample versus 7.5 ppts in the strongly balanced panel.⁸ Hence, most of the reallocation we document earlier occurs 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 average early in the sample experienced a relative fall

⁷If 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 years from now, it looks indistinguishable from that of the strongly balanced sample.

⁸This 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 their labor share afterwards, while smaller establishments instead saw a rise. We take 1982 as the starting point, which coincides with the start of the manufacturing labor share decline.⁹ In the context of the decomposition in Equation (2), this scenario would be akin to a negative correlation between size in 1982 and labor share changes over the subsequent 1982–2012 period. As a result, the covariance between labor shares and market shares, $Cov(\lambda_{it}, \omega_{it})$, would decline over time. We refer to this scenario as the big-player economy; it could be driven, for example, by the ability of larger establishments to lower wages through increasing monopsony power.

Alternatively, establishments with a low initial labor share may have grown faster than their peers in the following decades, a scenario we label the superstar economy. In this case, the decreasing covariance between labor shares and market shares would be driven by the fact that market share growth over the 1982–2012 period is negatively correlated with labor shares in 1982 (i.e., $Cov(\lambda_{i1982}, \Delta\omega_{it}) < 0$).

Both are examples of the composition effects we discuss in Section 3.1 and are 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 manufacturing labor shares. First, we compute the manufacturing labor share by keeping an establishment’s value-added weight ω_{it} equal to its 1982 value 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

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

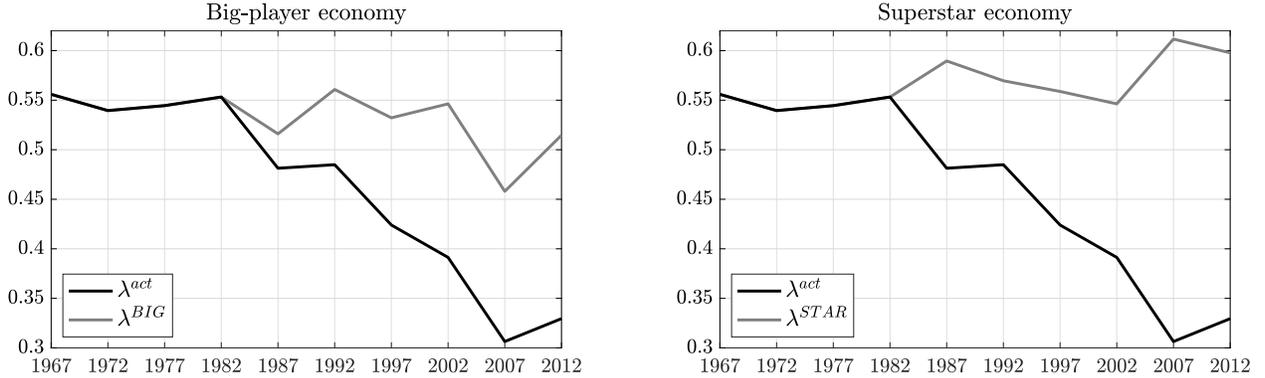
If either λ_t^{BIG} or λ_t^{STAR} were to coincide with the actual manufacturing labor share, we would conclude that either scenario deserves further attention. To conduct this analysis, we focus on a strongly balanced sample of establishments that are permanently active from 1982 to 2012. We display the actual and counterfactual labor shares in Figure 5.

The main takeaway from Figure 5 is clear: none of the two counterfactuals can reproduce the historical 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 remains more or less stagnant, while it dropped by about 18 ppts in the data over the same time period. In other words, the fall in the manufacturing labor share does not appear to be driven by an initial divergence in the relative labor shares of large versus small low labor share establishments.

⁹As a robustness check, we also consider 1977 as a starting point or as defining the “initial values” as the average around the 1982 Census: $\bar{\omega}_{i1982} = E_i[\omega_{i\tau}], \tau = 1977, 1982, 1987$.

¹⁰We thank an anonymous referee for suggesting these scenarios.

Figure 5: The Limited Role of Establishments with an Initially Small Size or Low Labor Share



Note: The figure plots the actual manufacturing labor share against the counterfactual labor shares in the big-player economy (left panel) and the superstar economy (right panel). It shows that neither did establishments that were initially relatively large experienced a particularly strong labor share decline. Nor did establishments with an initially relatively low level of the labor share grew their value-added shares.

If we instead freeze the labor shares at their 1982 level and focus on the contribution of the value-added weights (the superstar economy λ_t^{STAR}), we obtain a noticeable *rise* in the counterfactual labor share: from 55 ppts in 1982, it increases to 60 ppts by 2012, in contrast to a drop of 33 ppts in the actual data. The takeaway from this is that 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.

The importance of joint dynamics in the full panel Finding 1 made clear that the downward trend in the manufacturing labor share was due to a strong decline in the establishment-level covariance between labor and market shares, $Cov(\lambda_{it}, \omega_{it})$. The conclusion from the two exercises in the previous section is that *neither* market share dynamics *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, $Cov(\Delta\lambda_{it}, \Delta\omega_{it})$, at the micro level are key to our understanding of the aggregate trends in $Cov(\lambda_{it}, \omega_{it})$ and the manufacturing labor share λ_t . This analysis, however, requires focusing on a balanced panel. To confirm the crucial role played by the joint dynamics, we next perform a shift-share decomposition at higher frequency and for the full panel, in the spirit of [Baily, Hulten, and Campbell \(1992\)](#):

$$\Delta\lambda_t = \underbrace{\sum_i \omega_{it-5} \Delta\lambda_{it}}_{\text{Shift: } +6.9 \text{ ppt}} + \underbrace{\sum_i \Delta\omega_{it} \lambda_{it-5}}_{\text{Share: } +5.5 \text{ ppt}} + \underbrace{\sum_i \Delta\omega_{it} \Delta\lambda_{it}}_{\text{Interaction: } -15.1 \text{ ppt}} . \quad (9)$$

The first term in Equation (9) is akin to the big-player scenario we described earlier, in which value-added weights are kept fixed to their values in the previous census. The second term in-

stead freezes the micro-level labor shares at their past value while allowing value-added weights to fluctuate, similar to the superstar economy counterfactual. The last term captures the joint dynamics, corresponding to the interaction between the changes in the labor share and changes in the value-added share.

The numbers that accompany Equation (9) confirm our earlier findings based on the balanced sample. Specifically, the first two terms do not contribute to the decline in the manufacturing labor share. In fact, they contribute *positively* to the evolution of the manufacturing labor share, across all sample periods considered: for example, the contribution of the change in the labor shares holding value-added weights constant is +6.9 ppts from census to census over the 1982–2012 period, while that of the change in the weights alone is +5.5 ppts. The dramatic decline in the manufacturing labor share therefore has to come from the interaction term. Indeed, it has contributed about –15.1 ppts from census to census since 1982, thus overcompensating the positive contributions of the “shift” and “share” terms.

The main takeaway from the findings in this section is that the decline in the manufacturing 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 fell as they grew in size. These joint dynamics lead to a strong reallocation of economic activity toward the left tail of the labor share distribution and cause $Cov(\lambda_{it}, \omega_{it})$ 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 describe 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, we would expect this to then generate higher markups, leading to a fall in those establishments’ labor shares λ_{it} and a rise in their economic weights ω_{it} . The same negative empirical relationship between λ_{it} and ω_{it} would follow from some establishments gaining monopsonistic power in labor markets.

Yet, distinguishing between the various scenarios that we analyze in Section 3 ultimately requires a deeper analysis of the micro-level dynamics of labor shares and value-added shares, which we turn our attention to for the rest of this paper. In the next section, we identify the respective roles of value added, employment and wages in driving fluctuations in establishments’ labor shares.

4.4 Finding 3: Labor Shares Are Driven by Value Added, Not Wages or Employment

The first two findings imply that the factors behind the decline in the manufacturing labor share must (1) catalyze a reallocation of economic activity toward low labor share establishments and (2) generate a negative correlation between labor share and value-added dynamics at the establishment level. As we saw in the context of the conceptual framework of Section 3, all three types of shocks we discussed—demand, technology or monopsony—are, under some assumptions, consistent with this evidence. To discriminate between them, we now turn our attention to the cross-sectional

and time-series properties of the components of the labor share: wages, value added, employment, product prices and quantities.

4.4.1 Wages and Labor Productivity across Establishments

As a first step, let us rewrite the log of the labor share of establishment i at time t as

$$\log \lambda_{it} = \log W_{it} - \log ARPL_{it}, \quad (10)$$

where W_{it} is the average employee's wage and $ARPL_{it} = P_{it}Y_{it}/L_{it}$ denotes the average revenue product of labor.¹¹ As we saw in the conceptual framework, monopsony power in labor markets would predominantly affect wages W_{it} , while theories of efficiency or demand factors would impact $ARPL_{it}$.

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 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 state and three-digit NAICS industry.¹² The relative wage, \tilde{w}_{it} , and revenue labor productivity, $\widetilde{p_{it}y_{it}/l_{it}}$, are then defined in logs as follows:

$$\tilde{x}_{it} \equiv \log X_{it} - \overline{\log X_{-i,t}} \quad \text{where } \overline{\log X_{-i,t}} \equiv \sum_{j \neq i} \frac{P_{jt}Y_{jt}}{\sum_{j \neq i} P_{jt}Y_{jt}} \log X_{jt} \quad \text{and } X_{it} = W_{it}, \frac{P_{it}Y_{it}}{L_{it}}, \quad (11)$$

where we omit the industry and region subscripts for expositional purposes. The measure \tilde{x}_{it} is, by definition, centered around zero and denotes 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 (\tilde{w} and $\widetilde{py/l}$) in the cross-section. To do so, we run the following nonparametric regressions:

$$\tilde{x}_{it} = f(\lambda_{it}) + \varepsilon_{it}, \quad \tilde{x}_{it} = \tilde{w}_{it}, \widetilde{p_{it}y_{it}/l_{it}}, \quad (12)$$

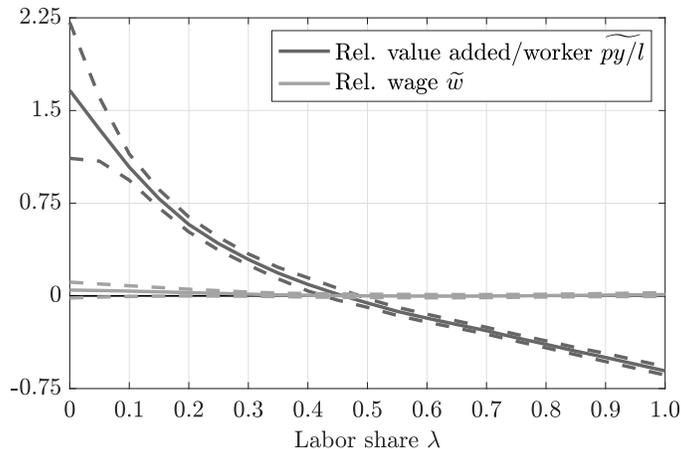
where \tilde{x}_{it} is either establishment i 's relative wage, \tilde{w}_{it} , or labor productivity, $\widetilde{p_{it}y_{it}/l_{it}}$. The function $f(\cdot)$ is the object of interest: it indicates whether low labor share establishments pay lower wages

¹¹It 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 differentiate between revenue labor productivity and physical labor productivity, the analog of physical total factor productivity (TFPQ) in Foster, Haltiwanger, and Syverson (2008); Hsieh and Klenow (2009).

¹²We 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 significantly change the conclusions.

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 value-added share (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 \tilde{w} and $\widetilde{py/l}$ are defined within each industry, year and region, we ensure that our findings are not driven by systematic differences along those dimensions.

Figure 6: 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 \tilde{w} against the labor share. All relative measures denote log-point differences vis-à-vis their peers as defined in Equation (11). Dashed lines denote 95% error bands.

The results of the two nonparametric regressions are displayed in Figure 6. They paint a clear picture. First, relative wages are nearly orthogonal to the labor share: establishments with a low labor share do not, on average, pay their workers more or less than their peers.¹³ 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 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. The takeaway from this analysis is that low labor share establishments do not

¹³Note that our estimate’s error bands denote the noise across establishments, not workers. Weighting observations (establishments) by their number of employees would reflect the more dispersed wage dispersion observed in worker- or household-level data. Even though we choose the more conservative establishment-level relative wage, the 95% error bands always include zero.

pay lower wages than their peers, as would be expected under theories of the labor share decline that rely on labor market power. Instead, they generate high value added per worker, which is consistent with theories of superior efficiency or consumer preferences.

4.4.2 Dynamics

In Section 4.3, we show that the joint dynamics of the labor share and value added at the establishment level are central to the aggregate behavior of the labor share. Next, we delve deeper into these dynamics by focusing on establishments at the bottom of the labor share distribution. This group, as we show in Section 4.2.2, experience a dramatic rise of its economic importance between 1967 and 2012. We start by defining and characterizing this subsample before studying its dynamics.

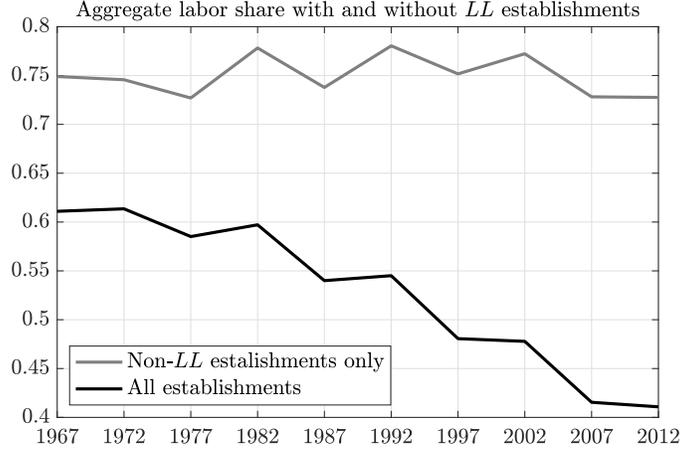
Defining low labor share establishments We define low-labor share (*LL*) establishments as those in the lowest quintile of the labor share distribution, in a given year and three-digit NAICS industry. The quintiles are industry specific due to the wide range of average labor shares across industries.

To highlight the role of *LL* establishments in shaping aggregate dynamics, we start by re-computing the manufacturing labor share without them. If reallocation toward lower labor shares was pervasive throughout the distribution in general, we would expect to also observe a labor share decline in the subsample without *LL* establishments, albeit from a higher starting point. The labor shares including and excluding *LL* establishments are shown in Figure 7. Two aspects stand out: first and unsurprisingly, 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, the level does not exhibit any 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 discernible downward trend in the second half of the sample. 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 labor share decline.

Going back to the cross-sectional analysis of Section 4.4.1, we find that *LL* establishments have an average relative labor productivity, \widehat{py}/l , 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 are not significantly different than their peers.

The dynamics of the labor share components Next, we investigate how the dynamics of the typical *LL* establishment’s labor share, and that of its components, differ from those of non-*LL* establishments. Using the data, our objective is to decompose the growth rate of micro-level labor shares ($\Delta \log \lambda_{it}$) into the contributions from wages ($\Delta \log W_{it}$), employment ($\Delta \log L_{it}$) and value

Figure 7: The Importance of *LL* Establishments for the Manufacturing Labor Share Decline



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 gray line). It shows that non-*LL* establishments do not contribute to the decline of the manufacturing labor share.

added ($\Delta \log(P_{it}Y_{it})$):

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

Our strategy is to use a regression approach to quantify the change of a specific variable for *LL* establishments *relative to their peers*. In particular, we first construct the growth rates in the labor share, employment, wage bill and value added between two census years (from years $t - 5$ to t) for each establishment in the 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. For example, for the growth rate of the labor share, the specification is

$$\Delta \log x_{it} = c + \beta \mathbb{I}\{LL_{it}\} + \gamma X_{it} + \varepsilon_{it} \quad \text{where } x_t = \lambda_{it}, W_{it}, L_{it} \text{ or } P_{it}Y_{it}. \quad (13)$$

While the *level* of the labor share of *LL* establishments is below that of their peers by definition—they consist of *LL* establishments in the lowest quintile in a given year and industry—our aim here is to uncover their relative *dynamics* from the estimates of the coefficient β in Equation (13). That is, we study how the dynamics of the labor share and its components for the typical *LL* establishments differ from those of non-*LL* establishments over the previous five-year window. Note that we do not require that *LL* establishments in period t also be *LL* establishments in $t - 5$. The vector X_{it} contains industry, region and year dummies as controls. We estimate Equation (13) with and without value-added weights to account for the fact that larger establishments are likely to have less volatile labor shares. The procedure is similar for the wage bill, employment and value added regressions.¹⁴ Results from the weighted regressions are displayed in Table 1.

¹⁴We also study the dynamics of capital intensity and intermediates and find little evidence that they are different

Table 1: Dynamics of LL Establishments

	(I)	(II)	(III)	(IV)
	$\Delta \log \lambda_{it}$	$\Delta \log W_{it}$	$\Delta \log L_{it}$	$\Delta \log(P_{it} Y_{it})$
β	-0.4632*** (0.0154)	-0.0099 (0.0100)	0.0001 (0.0284)	0.4532*** (0.0442)
Change in ppt	-18.04	-0.4	0.0	-17.7
R^2	0.186	0.135	0.021	0.114

Note: This table shows the pooled OLS regression of Equation (13) on the full census panel. Observations are weighted using the share of establishment i 's value added in overall manufacturing value added. Standard errors are clustered at the four-digit NAICS industry level. Significance levels are denoted by * (10% level), ** (5% level) and *** (1% level).

The first column of Table 1 implies that relative to the previous census year, an establishment that has LL status at time t saw its labor share fall by 46%. This strongly significant estimate translates into a labor share drop of 18 ppts. Columns (II)–(IV) of Table 1 present the results from a similar value-weighted regression but with the relative wage, employment or value added on the left-hand side of the equation. They indicate that out of the 18 ppt drop in the labor share for the typical LL establishment, a full 17.7 ppts come from increasing value added relative to non- LL establishments. In contrast, the relative dynamics of wages and employment do not contribute to the differential labor share dynamics of LL establishments in a meaningful way. Note that when we estimate the relative dynamics in an unweighted fashion, the results are even stronger, suggesting a relative labor share decline of 29 ppts, on average, for LL establishments. Again, an overwhelming proportion of this change is driven by value added.

Going back to the framework of Section 3.2, recall that we could write the labor share as $\lambda_{it} = \frac{\nu_{it} MRPL_{it}}{ARPL_{it}}$. Since we find no evidence that wages are important for explaining labor share differences across establishments or over time, we can conclude that the “exploitation parameter,” ν_{it} , is not a quantitatively important factor driving labor shares in our sample. In sum, it appears unlikely that increased monopsony power in the labor market is behind the fall in the manufacturing labor share.

4.5 Finding 4: Low Labor Shares Stem Mostly from a Product Price Premium

The previous section highlights the key role played by value added: cross-sectional and dynamic differences between LL and non- LL establishments 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 provide evidence that demand-side factors, rather than technology, appear to be a key driver of micro-level labor share patterns.

for LL establishments relative to their peers.

4.5.1 Measuring Prices

To identify the relative contributions of these two distinct forces, we turn to another data source provided by the U.S. Census Bureau: the product trailer to the CMF. 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 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 ten-digit NAICS level, a finer definition of product than most of the literature.¹⁵ This is a product-coding system devised by the Census Bureau and is 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](#)).¹⁶

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 census. 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 pounds), 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 130,000 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 find to be the key driver of labor shares in the cross-section and time series. Since all price data are sales based, we switch to studying sales per worker, rather than value added per worker, when analyzing the price versus physical productivity difference. We define relative sales per worker analogous to that of relative value added per worker

¹⁵[Foster, Haltiwanger, and Syverson \(2008\)](#) define products at the seven-digit SIC code level, while [Bernard, Redding, and Schott \(2010, 2011\)](#) aggregate product sales to the five-digit SIC level of products; both definitions are coarser than ours.

¹⁶[White, Reiter, and Petrin \(2018\)](#) show that the product trailer data set is seriously contaminated by imputations based on industry averages or regression models.

in Equation (11):

$$\widetilde{p_{it}q_{it}/l_{it}} \equiv \log(P_{it}Q_{it}/L_{it}) - \overline{\log(P_{-i,t}Q_{-i,t}/L_{-i,t})} \quad (14)$$

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

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

4.5.2 Product Prices across Establishments

To make prices comparable across establishments, we adopt the treatment of nominal wages and labor productivity in Section 4.4 by comparing establishment-level prices to a peer group. This time, however, we have to start at the product level. First, we normalize prices at the level of the ten-digit NAICS product ℓ :

$$\widetilde{p_{i\ell t}} \equiv \log P_{i\ell t} - \overline{\log P_{-i,\ell t}} \quad \text{where } \overline{\log P_{-i,\ell t}} \equiv \sum_{j \neq i} \frac{P_{j\ell t}Q_{j\ell t}}{\sum_{j \neq i} P_{j\ell t}Q_{j\ell t}} \log P_{j\ell t}. \quad (15)$$

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} :

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

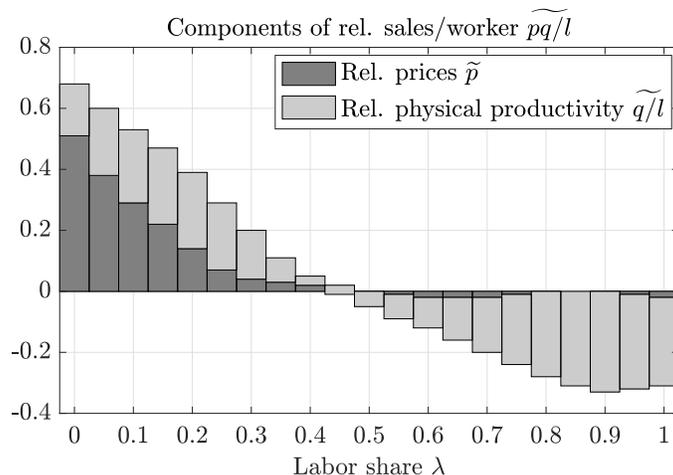
We refer to \widetilde{p}_{it} as the average product price premium that establishment i charges relative to its peers across its product lines. This measure represents the mean log-point difference between an establishment's output prices and those of its peers.¹⁷

Similar to our earlier approach, we nonparametrically 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

¹⁷A 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.

relative sales per worker explained by prices; the remainder is the portion explained by physical labor productivity $\widetilde{q/l}$. If establishments with a low labor share operated superior technologies and produced the same goods ℓ more efficiently, we would expect them to post lower prices in those categories and to sweep up the market at the expense of their peers. Such technology-driven growth would show up as a generally negative \widetilde{p} for establishments with a low labor share and vice versa for high labor share establishments. If, on the other hand, low labor share establishments faced favorable demand conditions that allowed them to post a higher price and to generate higher revenues as a result, we would anticipate an opposite pattern: positive \widetilde{p} for low labor share and negative \widetilde{p} for high labor share establishments.

Figure 8: The Contributions of Physical Productivity and Prices to Relative Sales per Worker



Note: The figure displays the cross-sectional differences in relative prices \widetilde{p} (dark gray bars) and relative physical labor productivity $\widetilde{q/l}$ (light gray bars) against the labor share; \widetilde{p} defined in Equation (15) and $\widetilde{q/l}$ is the ratio of $\widetilde{pq/l}$ (defined in Equation (14)) and \widetilde{p} .

The contributions of the two components (relative prices and relative physical productivity) to differences in relative sales are depicted in Figure 8. We can see that while relative sales per worker are driven by both price and physical labor productivity differences, prices 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 establishments' sales/workers with a labor share of 50 pts 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 $\exp(0.430 + 0.096) \approx 69\%$.

4.5.3 Dynamic Evidence

Analogous to the dynamic analysis of wages, employment and value added in the previous section, we repeat the estimation of Equation (13) for relative prices, $\Delta\tilde{p}_{it}$, and relative physical output per worker, $\Delta(\widetilde{q_{it}/l_{it}})$. The results are displayed in Table 2.

Table 2: Price and Physical Productivity Dynamics of *LL* Establishments

	(I)	(II)
	$\Delta\tilde{p}_{it}$	$\Delta(\widetilde{q_{it}/l_{it}})$
β	0.1675*** (0.0545)	-0.0493 (0.0844)
R^2	0.093	0.135

Note: This table shows the pooled OLS regression of Equation (13) on the matched price panel where each observation is weighted with its share in manufacturing sales. For more details, see notes to Table 1.

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). While the coefficient estimates for the *LL* versus non-*LL* price dynamics are statistically significant, this is not the case for physical productivity: the point estimate, at -4.9%, is more than three times smaller than for prices and is not significantly different than zero.

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 Findings 1 to 4: preference or TFP shocks, combined with non-isoelastic demand schedules, can explain the joint dynamics of labor and market shares at the establishment level because both of them would increase markups, leading to a reallocation of economic activity toward *LL* establishments. Yet, the fact that relative prices and labor shares co-move negatively represents strong evidence in favor of demand shocks suggesting that demand shocks are first order to explain labor share dynamics of *LL*: under technology shocks, we would expect relative prices to *fall* alongside labor shares. Furthermore, we provide additional evidence on the unimportance of supply factors in Appendix D.

4.6 Finding 5: Shooting Stars, or the Transience of Low Labor Shares

Seen through the lens of the conceptual framework of Section 3.2, our evidence indicates that demand factors must be playing a central role: they can rationalize both the joint dynamics of labor share and value added as well as the importance of prices in driving the high nominal labor productivity of low labor share establishments. The analysis of *LL* establishments also show that this status is the product of an economically large rise in value added, driven mainly by higher prices. But what is the dynamic nature of these underlying demand drivers? Is their impact on the labor shares of establishments highly persistent or transient? The answer is relevant at many

levels. For one, it can instruct policymakers on the nature of concentration: transient labor and market shares would have different implications for competition policy than if the economy was characterized by *LL* establishments that are progressively taking over their market and are lowering their labor shares. Moreover, it can help us have a better sense of the nature of demand factors and their impact on firms' actions. With these objectives in mind, in this section, we turn to an analysis of the labor share persistence at the micro level.

Markov transitional dynamics We start by analyzing the transition dynamics of *LL* and non-*LL* establishments. Our objective is to assess whether the demand drivers identified in Findings 3 and 4 are important enough to perturb the rankings of establishments along the labor share dimension. We do so with the help of a Markov transition matrix, displayed in Table 3. More specifically, 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 *LL* status was highly persistent, this probability should be equal to 100%. At the polar opposite, if labor shares are so volatile that they perturb the ranking every period, we should expect the identity of *LL* establishments to be random and the transition probability to be close to 20%.

Table 3: Transition Probabilities of *LL* Status

Panel A: Unweighted transitional dynamics		
	Non- <i>LL</i> _{$t+5$}	<i>LL</i> _{$t+5$}
Non- <i>LL</i> _{t}	0.854	0.146
<i>LL</i> _{t}	0.583	0.417
Panel B: Weighted transitional dynamics		
	Non- <i>LL</i> _{$t+5$}	<i>LL</i> _{$t+5$}
Non- <i>LL</i> _{t}	0.922	0.078
<i>LL</i> _{t}	0.536	0.464

Note: The table shows the 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, and Panel B displays the share of manufacturing value added accounted for by the *LL* establishments when defined by *VA*-weighted quintiles.

Table 3 shows that over our sample period, the probability that an establishment retains *LL* status from census year to census year (a five-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 3 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; while they indicate slightly more persistence, the overall impression remains unchanged.

V-shaped dynamics In light of their surprisingly temporary nature, we aim next to quantify the labor share dynamics that occurs in the years following *LL* status. To do so, we adopt the same type of regression framework as in Sections 4.4.2 and 4.5.3, which captures the dynamics of *LL* establishments relative to their peers. Specifically, we regress both backward-looking (from years $t-5$ to t) and forward-looking (from t to $t+5$) percentage point changes in establishment-level labor shares on a dummy variable that equals one if an establishment is among the *LL* establishments in the current census year:

$$\Delta \log \lambda_{it} = c_1 + \beta_{-5} \mathbb{I}\{LL_{it}\} + \gamma X_{it} + \varepsilon_{1it} \quad (16)$$

$$\Delta \log \lambda_{it+5} = c_2 + \beta_{+5} \mathbb{I}\{LL_{it}\} + \gamma X_{it} + \varepsilon_{2it}. \quad (17)$$

Our objective is to rely on the estimates of the coefficients β_{-5} and β_{+5} in Equations (16) and (17) to characterize 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 also be *LL* establishments in $t-5$ and in $t+5$; an establishment could well have *LL* status for a single year. The vector X_{it} contains controls that have been described earlier, and estimation is done both with and without value-added weights. Results are displayed in Table 4. Since Equations (13) for λ_{it} and (16) here are equivalent, the weighted results in the first row here are exactly the dynamics shown in Column (I) of Table 1.

Table 4: Before-After Dynamics of *LL* Establishments

	(I)	(II)	(III)	(IV)	(V)	(VI)
Variable	$\Delta \log \lambda_t$	$\Delta \log \lambda_{t+5}$	$\Delta \lambda_t$	$\Delta \lambda_{t+5}$	$\Delta \lambda_t$	$\Delta \lambda_{t+5}$
β_{-5}	-0.4632*** (0.0154)		-0.1804*** (0.0100)		-0.2896*** (0.0076)	
β_{+5}		+0.3844*** (0.0177)		+0.1542*** (0.0102)		+0.2710*** (0.0071)
R^2	0.186	0.108	0.102	0.070	0.111	0.096
VA weights	yes	yes	yes	yes	no	no

Note: The table shows the pooled OLS regression of Equations (16) and (17) on the full census panel. For more details, see notes to Table 1.

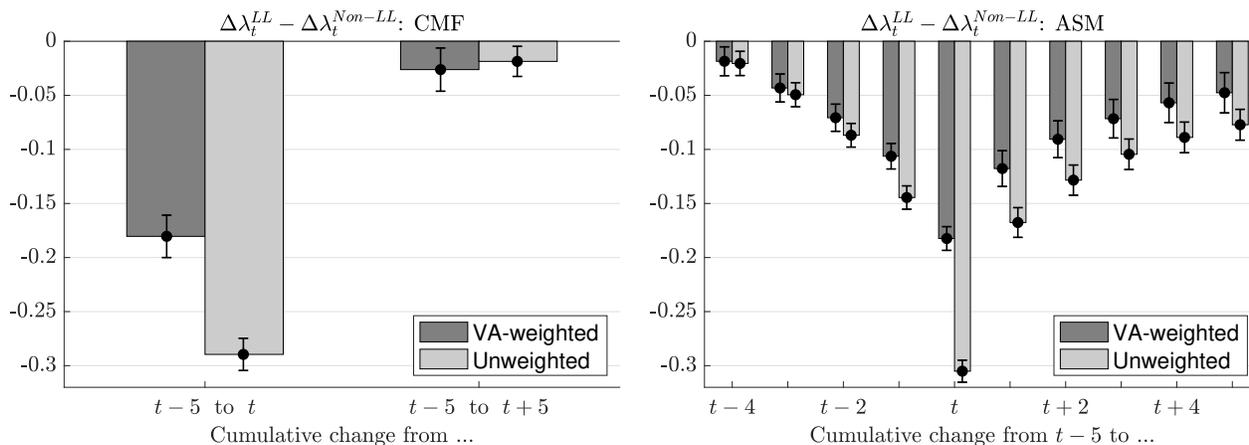
If the labor shares of *LL* establishments were to permanently reach a lower level relative to their non-*LL* peers, the coefficient β_{+5} would be close to zero. The estimation indicates, however, that this is clearly not the case: β_{+5} is statistically different than zero. In fact, it shows again that *LL* status is highly transient: while a typical establishment with *LL* status at time t saw its labor share since $t-5$ shrink by 46.3% (Column (I)), it *rises* by 38.4% in the subsequent five years (Column (II)).

To ease interpretation, we repeat the regressions (16) and (17) with the percentage point change of labor shares, $\Delta \lambda_{it} = \lambda_{it} - \lambda_{it-5}$, on the left-hand side instead of the growth rate. As Column

(III) shows, *LL* establishments experience a relative labor share decline of 18 ppts. Yet, in the five-year period thereafter (from t to $t + 5$), the coefficient estimates of β_{+5} in Columns (II)/(IV) indicate that the change in the labor share of establishments that are *LL* in year t will expand by 38.4% (corresponding to 15.4 ppts) *more* than that of non-*LL* establishments.

Finally, we report the results for β_{-5} and β_{+5} as *cumulative* growth rates in the left panel of Figure 9. The figure confirms that compared to $t - 5$, time- t *LL* establishments appear to be barely different than their non-*LL* peers by $t + 5$. The unweighted estimates in Columns (V) and (VI) are stronger, indicating that small *LL* establishments see even more extreme labor share dynamics in absolute value. All in all, our analysis appears to show that the average *LL* establishment experiences a rather temporary drop and rebound in its labor share. This finding is in line with the earlier evidence from the Markov transition matrices and confirms the low persistence of *LL* status.

Figure 9: The Temporary Fall and Rise of Labor Shares of *LL* Establishments



Note: The left panel shows the 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 are in dark gray, value added-weighted dynamics are in light gray, and whiskers denote 95% error bands. The y-axis represents labor share changes, where the labor share is expressed as a decimal. The right panel shows analogous labor share dynamics of *LL* establishments in the Annual Survey of Manufactures 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 9. 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 *LL* establishments, measurement error for large establishments whose labor share is low is much less likely, as the Census Bureau 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 data set does not capture the population of manufacturing establishments, its macro-level 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 have LL status 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 (16) and (17) 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 9. They confirm the transient nature of the labor share that we found using 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 occurs 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 5–8 ppts 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.¹⁸

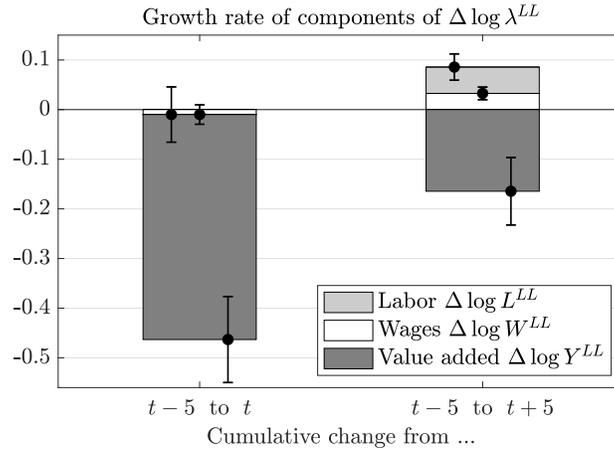
4.6.1 Why Does the Labor Share Rebound?

In Section 4.3, we find that the drop in the labor share of the typical LL establishment was due to a strong increase in value added. In the context of the V-shape dynamics discussed earlier, this naturally leads to another question: is the rebound of the LL labor share in the following five years driven by employment or wages catching up with revenue labor productivity? The former would suggest that demand shocks, while large and dominant in driving labor shares, are rather transient; the latter, on the other hand, could result from labor adjustment costs or rigid wages. To quantify the relative contributions of wages, employment and value added to the t to $t + 5$ labor share dynamics of LL establishments, we estimate Equations (16) and (17) for these three variables separately. The results are displayed in Figure 10.

The leftmost bar in Figure 10 depicts graphically the results from Table 1: between $t - 5$ and t , the average LL establishment saw its labor share shrink by 18 ppts compared to the typical non- LL establishment, and this drop was almost entirely due to the differential in value added growth. The rightmost bar instead shows the *cumulative* contribution of each component. Incorporating the five following years, we see that the V-shaped 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

¹⁸As 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 an establishment’s individual products. 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 9. In that exercise, we drop sales recorded under product balancing codes and omit imputed product values to guard against the problems associated with imputation highlighted in White, Reiter, and Petrin (2018), so the sum of product sales and total value of shipments do not necessarily coincide.

Figure 10: 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) and 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.

retreat of value-added growth accounts for $11 \frac{1}{2}$ ppts of the $15 \frac{1}{2}$ ppt rebound in the average *LL* labor share, whereas wages and employment do pick up but contribute only 2 ppts each to the labor share rebound. In other words, there is little contribution coming from a delayed response of employment and wage growth.

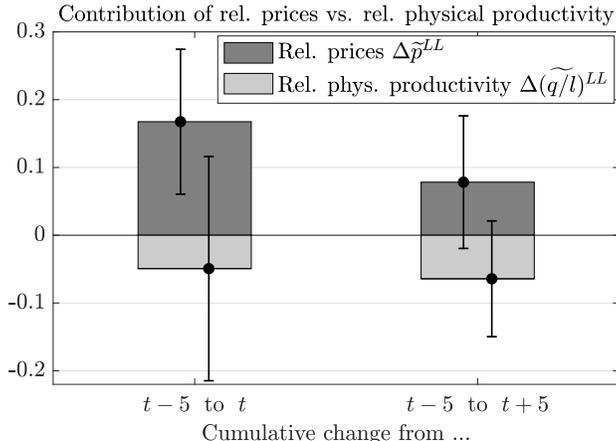
4.6.2 Why Does Value Added Recede?

Using the matched price sample, we showed in Section 4.5 that the increase in nominal sales per worker of *LL* establishments was a product of higher relative prices, not higher physical productivity. Using a similar regression approach, we also find that in the subsequent five years most, but not all, of that jump in the product price premium is reverted. As can be seen in Figure 11, the cumulative change in average product price premium from $t - 5$ to $t + 5$ is only 7.8% higher for those establishments that are *LL* at time t relative to their peers, while it is 16.8% between $t - 5$ and t ; only the initial 16.8 ppt jump is statistically significant. The dynamics of physical productivity over the two horizons are, by contrast, similar in magnitude and are not statistically significant.

5 Did Demand Shocks Become More Important Over Time?

The previous section distilled a number of empirical findings that identified the key drivers of the manufacturing labor share decline. First, the increasingly negative co-movement of labor shares and value-added shares at the establishment level was crucial for the decline from the 1980s onward (Findings 1 and 2). Second, using our conceptual framework from Section 3, we provided additional empirical facts pointing to demand factors as the engine of these aforementioned micro-level

Figure 11: Relative Price Dynamics Dominate Relative Sales Dynamics



Note: The figure displays the dynamic contributions of the growth in relative prices $\Delta \tilde{p}$ and labor productivity growth $\Delta(\tilde{q}/l)$ for sales per worker growth $\tilde{p}\tilde{q}/l$ 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.

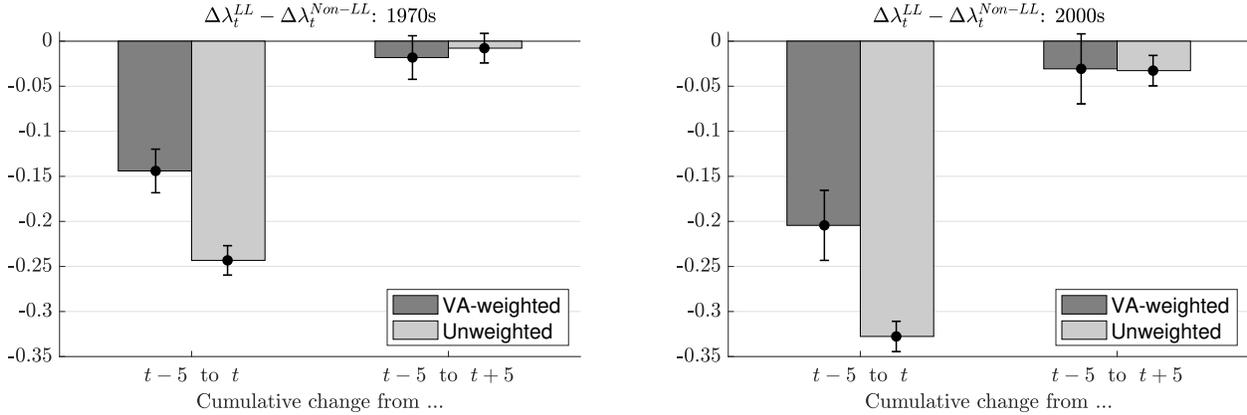
dynamics (Findings 3–5). Can these demand factors be behind not only the micro-level dynamics but also the decline of the labor share in the manufacturing sector? In this section, we provide evidence of significant changes in the micro-level anatomy of labor shares and their components over our sample period. In the context of the conceptual framework of Section 3.2, we contend that these findings are consistent with a rise in the volatility of demand-side factors.

5.1 The V-Shaped Labor Share Pattern Gets Deeper Over Time

We start by investigating the evolution over time of the V-shaped labor share pattern of LL establishments that we documented earlier. We repeat the dynamic analysis described by Equations (16) and (17) separately for the 1972 and 1977 censuses, denoted as “1970s,” and the 2007 and 2012 censuses, denoted as “2000s.” Both unweighted and value-added weighted estimates for these equations are shown in Figure 12.

Focusing first on the $t-5$ to t dynamics, we find that the labor share dynamics of LL and non- LL establishments get increasingly different over time. In the weighted case, comparing the left and right panels indicates that the differential increased by 50%, from a relative 14 ppts, on average, in the 1970s to a relative 21 ppts by the 2000s. Our earlier finding that labor share dynamics are very transient, on the other hand, appears to hold over time: between $t-5$ and $t+5$, the (weighted) cumulative differential is about -2 ppts in the 1970s versus -3 ppts in the 2000s, with neither estimate being statistically different from zero. In the unweighted case, the $t-5$ to t differential is 24 ppts, on average, in the 1970s and increases to 33 ppts by the 2000s. Again, most of this difference disappears once we consider a ten-year window centered around year t .

Figure 12: Labor Share Change of LL versus Non- LL Establishments Over Time

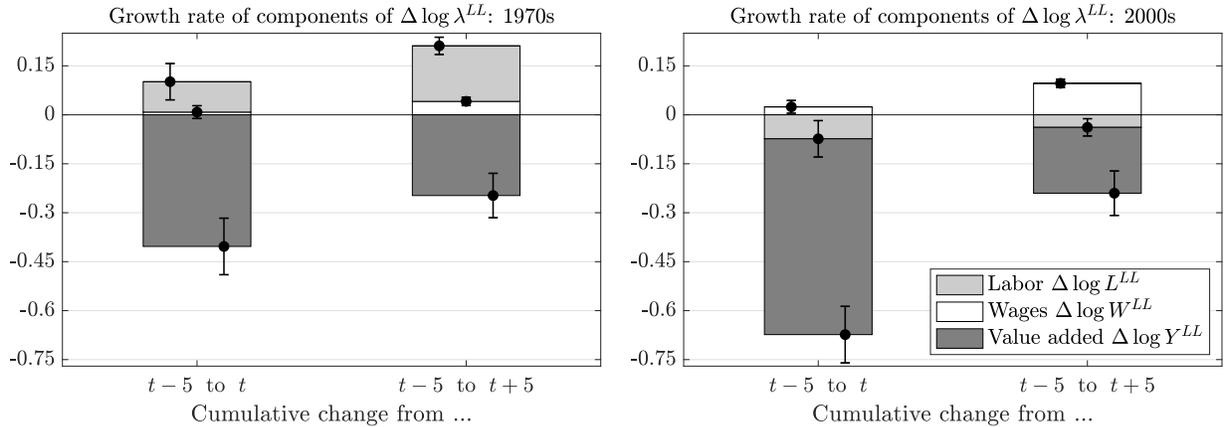


Note: This figure displays the difference in labor share dynamics between LL and non- LL establishments (corresponding to the $t - 5$ to t bars in the left panel of Figure 9) by year. It shows that LL establishments look strongly different from their peers in the 2000s.

5.2 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. In fact, we show next that the co-movement 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). We repeat the exercise of Section 4.6.1 by estimating Equations (16) and (17) for wages, employment and value added but this time for the early and late samples separately. Figure 13 displays the contributions of these three components to the labor share growth rate of LL establishments relative to non- LL establishments.

Figure 13: 1970s versus 2000s



In the 1970s, the majority of the adjustment in the five-year period preceding census year t was driven by a rise in value added (negative contribution to the labor share): the average LL establishment’s labor productivity growth was 40 ppts 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 ppts higher for LL establishments. In the five following years, almost all the labor share growth differential disappears. This occurs 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 does so with a delay. This picture is consistent with the notion that hiring frictions delay the employment response of LL establishments, which tie it instead to the longer run dynamics of value added. Ultimately, while the value added of LL establishments has clearly grown more over the ten-year span than that of their peers, relative labor productivity is more or less back to where it was initially because employment and, to a lesser extent, wages catch up with value added.

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 60 ppts instead of 40 ppts in the 1970s. Second, the V-shaped pattern is now more pronounced: not only is the value-added growth differential sharper initially, but after ten years, only 20 ppts remain in the 2000s, compared to 25 ppts in the 1970s. Third, the response of employment is noticeably different from the early part of the sample: between $t - 5$ and t , employment growth is 7 ppts *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 close to zero.

Taken together, the findings in this section highlight two significant developments in the dynamics of labor shares at the manufacturing establishment level. First, there has been a deepening over time of the V-shaped labor share pattern of LL establishments that we document in Section 4.6. This is due in large part to a sharper response of value added relative to their peers. In the context of the conceptual framework, this can be interpreted as an increase in the volatility of the demand factors that underlie the micro-level dynamics of the labor share. With such extremely positive demand shocks, LL establishments will find themselves in a very inelastic part of their demand curve, where most of the demand shock is passed through into higher prices rather than into higher employment. This means that our second documented change, the disconnect between value added and labor input, has become stronger over time. This is in line with recent work documenting the decline in the economy’s responsiveness to shocks; see, for example Table 5 in Ilut, Kehrig, and Schneider (2014) or the work by Pugsley, Sedláček, and Sterk (2017); Decker, Haltiwanger, Jarmin, and Miranda (2017a,b); Cooper, Haltiwanger, and Willis (2017).

5.3 Discussion

In Section 4, we argued that our evidence points toward demand as the main driver of establishment-level labor share dynamics. Hence, our finding above, that these dynamics have become more pronounced over time, would hint at demand factors being more volatile recently than they were

a few decades ago. But what could be behind these changes? In this section, we offer several potential explanations for this trend.

A number of developments in the economy have made the customer increasingly important for businesses. First, increasing market integration has raised variety. This has allowed customers to choose from a larger set of products and firms, at a wider geographical scale. As a result of increased market reach, the gains from being the leading firm have become larger and increasingly profitable. In terms of our conceptual framework, this “winner-take-all” environment is akin to more volatile demand shocks.

These potential gains may arguably lead firms to expand additional resources to sell the product that will be highly sought after in the future or to make their customer base more immune to competition. This could include more intensive advertising activity or selling new valuable services along with a product. Ultimately, by making their demand curve less elastic, demand shocks would translate into stronger price increases relative to the physical output and employment responses. In Appendix D.1, we provide evidence that *LL* establishments spend more on advertising than their peers and its effectiveness seems to have increased over time.

This higher volatility of demand shocks would have a number of implications in line with the findings we document earlier. First, we would expect an increasingly negative co-movement between value-added shares and labor shares (Finding 2), accompanied by more pronounced V-shaped labor share dynamics of *LL* establishments (Finding 5, Figure 12). Second, an increasingly prominent role of demand shocks would mean that firms respond to these shocks less and less by expanding employment while nominal value added is predominantly carried by price increases (Finding 4). This increasing non-responsiveness of employment to demand shocks results in a weakening of the co-movement between employment and value added (Finding 5, Figure 13). This increasing disconnect between profitability shocks that move nominal demand on the one hand, but with a merely lackluster response of employment on the other, is noted in the literature.

The increasing volatility of demand shocks raises some interesting questions about where these demand shocks occur. Are they specific to establishments (i.e., the units’ manufacturing specific goods are in high demand)? In other words, do *LL* establishments within the same firm undergo a V-shaped labor share pattern at the same time? In Appendix C, we show that the V-shaped patterns, their anatomy and their increasing salience are also present at the firm level. While a more thorough analysis of this dimension is beyond the scope of this paper, this evidence has significant implications. For one, it confirms that our findings are not just the byproduct of measurement error at the establishment level. More importantly, it indicates that the demand factors driving labor share dynamics more likely occur at the firm level. They could stem, for example, from brand power or from the propagation of demand shocks through the vertically integrated structure of the firm.

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. We find that this apparent disconnect is due to a drastic reallocation of economic activity toward the lower end of the labor share distribution. We show that this was not driven by entry/exit or compositional forces but instead by establishments whose labor share fell as they grew in size.

We then highlight a number of additional micro-level empirical facts. In particular, we show that low labor-share (*LL*) establishments are characterized by high revenue labor productivity, not low wages, and charge higher prices than their peers for similar products. Moreover, we find that *LL* status is very transient, a pattern that has become more pronounced over time. In the context of a simple conceptual framework, we conclude that among the leading theories proposed in the literature to explain the decline in the manufacturing labor share, only demand factors are consistent with all of our empirical findings.

These findings, taken as a whole, provide a guide for researchers intent on understanding and modeling 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 4.)
- Francisco Alvarez-Cuadrado, Ngo Van Long, and Markus Poschke. Capital-labor substitution, structural change and the labor income share. *Journal of Economic Dynamics and Control*, 87: 206–231, February 2018. (Cited on pages 4 and 8.)
- 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 2017. (Cited on pages 5 and 8.)
- 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. *Quarterly Journal of Economics*, forthcoming. (Cited on pages 5, 8, 14, and 46.)
- José Azar, Ioanna Marinescu, and Marshall I. Steinbaum. Labor market concentration. *NBER Working Paper No. 24147*, 2017. (Cited on page 9.)
- 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 17.)

- David Rezza Baqaee and Emmanuel Farhi. Productivity and misallocation in general equilibrium. *Quarterly Journal of Economics*, 135(1):105–163, February 2020. (Cited on pages 5 and 8.)
- Simcha Barkai. Declining labor and capital shares. *Working Paper*, 2017. (Cited on pages 5 and 8.)
- David Berger, Kyle Herkenhoff, and Simon Mongey. Labor market power. *NBER WP No. 25719*, March 2019. (Cited on page 9.)
- Daniel Berkowitz, Hong May, and Shuichiro Nishioka. Does capital-labor substitution or do institutions explain declining labor shares? *Working Paper*, 2017. (Cited on page 14.)
- 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 24.)
- 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 24.)
- 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 9.)
- 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 4.)
- Christoph E. Boehm, Aaron Flaaen, and Nitya Pandalai-Nayar. Multinationals, offshoring, and the decline of U.S. manufacturing. *Working Paper*, 2015. (Cited on page 4.)
- Benjamin Bridgman. Is labor’s loss capital’s gain? gross versus net labor shares. *Working Paper*, 2014. (Cited on page 5.)
- Erik Brynjolfsson, Andrew McAfee, Michael Sorell, and Feng Zhu. Scale without mass: Business process replication and industry dynamics. *Harvard Business School Technology & Operations Mgt. Unit Research Paper No. 07-016*, September 2008. (Cited on page 5.)
- Russell W. Cooper, John C. Haltiwanger, and Jonathan L. Willis. Declining dynamism at the establishment level. *SED Meeting Paper*, 2017. (Cited on page 35.)
- Steven J. Davis, John C. Haltiwanger, and Scott Schuh. *Job Creation and Destruction*. MIT Press, Cambridge, MA, 1996. (Cited on page 42.)
- Jan De Loecker, Jan Eeckhout, and Gabriel Unger. The rise of market power and the macroeconomic implications. *Working Paper*, November 2018. (Cited on pages 5 and 8.)
- Ryan A. Decker, John Haltiwanger, Ron S. Jarmin, and Javier Miranda. Declining dynamism, allocative efficiency, and the productivity slowdown. *American Economic Review Papers and Proceedings*, 107(5):322–326, May 2017a. (Cited on page 35.)
- 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 35.)
- 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 pages 42 and 52.)

- 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 43.)
- Sebastian Dyrda and Benjamin Pugsley. Taxes, private equity, and evolution of income inequality in the US. *Working Paper*, 2019. (Cited on page 59.)
- Maya Eden and Paul Gaggl. On the welfare implications of automation. *Review of Economic Dynamics*, 29:15–43, July 2018. (Cited on pages 4 and 10.)
- Chris Edmond, Virgiliu Midrigan, and Daniel Y. Xu. How costly are markups? *NBER Working Paper No. 24800*, 2018. (Cited on pages 5, 8, and 25.)
- Andrea Eisfeldt, Antonio Falato, and Mindy X. Zhang. Human capitalists. *Working Paper*, 2018. (Cited on page 44.)
- Michael W. L. Elsby, Bart Hobijn, and Aysegül Şahin. The decline of the U.S. labor share. *Brookings Papers on Economic Activity*, pages 1–63, Fall 2013. (Cited on pages 4, 5, 9, and 10.)
- 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 9.)
- 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 43.)
- 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 15.)
- 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 19, 24, and 47.)
- Jason Furman and Peter Orszag. A firm-level perspective on the role of rents in the rise in inequality. *Working Paper*, 2015. (Cited on page 4.)
- Andrew Glover and Jacob Short. Demographic origins of the decline in labor's share. *Working Paper*, 2018. (Cited on pages 4 and 9.)
- Emilien Gouin-Bonenfant. Productivity dispersion, between-firm competition and the labor share. *Working Paper*, 2018. (Cited on pages 9 and 14.)
- François Gourio and Leena Rudanko. Customer capital. *Review of Economic Studies*, 81(3):1102–1136, July 2014. (Cited on page 57.)
- 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 4 and 8.)
- Gustavo Grullon, Yelena Larkin, and Roni Michaely. Are US industries becoming more concentrated? *Working Paper*, 2016. (Cited on pages 5 and 8.)
- Germán Gutiérrez and Thomas Philippon. Declining competition and investment in the U.S. *NBER Working Paper No. 23583*, 2017. (Cited on page 8.)

- 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 4.)
- Brad Hershbein, Claudia Macaluso, and Chen Yeh. Labor market concentration and the demand for skills. *Working Paper*, 2018. (Cited on page 9.)
- Susan N. Houseman. Understanding the decline of U.S. manufacturing employment. *Upjohn Institute Working Paper 18-287*, January 2018. (Cited on page 43.)
- 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 19.)
- 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 35.)
- Gregor Jarosch, Jan Sebastian Nimczik, and Isaac Sorkin. Granular search, market structure, and wages. *NBER WP No. 26239*, September 2019. (Cited on page 9.)
- Nicholas Kaldor. Capital accumulation and economic growth. In F. A. Lutz and D. C. Hague, editors, *Theory of Capital*, pages 177–222. St. Martin’s Press, New York, 1961. (Cited on page 2.)
- Loukas Karabarbounis and Brent Neiman. The global decline of the labor share. *Quarterly Journal of Economics*, 129(1):61–103, February 2014a. (Cited on pages 4 and 8.)
- Loukas Karabarbounis and Brent Neiman. Capital depreciation and labor shares around the world: Measurement and implications. *Working Paper*, 2014b. (Cited on page 5.)
- Bariş Kaymak and Immo Schott. Corporate tax cuts and the decline of the labor share. *Working Paper*, 2018. (Cited on pages 4 and 8.)
- Matthias Kehrig. The cyclical nature of productivity dispersion. *US Census Bureau Center for Economic Studies Paper No. CES-WP-11-15*, May 2011. (Cited on pages 5 and 42.)
- Matthias Kehrig. The cyclical nature of the productivity distribution. *Working Paper*, 2015. (Cited on page 46.)
- Miles S. Kimball. The quantitative analytics of the basic neomonetarist model. *Journal of Money, Credit and Banking*, 27:1241–1277, 1995. (Cited on page 8.)
- Dongya Koh, Raül Santaèulàlia-Llopis, and Yu Zheng. Labor share decline and intellectual property products capital. *Working Paper*, 2019. (Cited on pages 4 and 10.)
- Robert Z. Lawrence. Recent declines in labor’s share in US income: A preliminary neoclassical account. *NBER Working Paper No. 21296*, 2015. (Cited on page 4.)
- Ash Leblebicioğlu and Ariel Weinberger. Credit and the labor share: Evidence from U.S. states. *Economic Journal*, forthcoming. (Cited on page 8.)
- 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 8.)
- Brent Neiman and Joseph Vavra. The rise of niche consumption. *NBER Working Paper No. 26134*, August 2019. (Cited on pages 5 and 8.)

- Ezra Oberfield and Devesh Raval. Micro data and macro technology. *NBER Working Paper No. 20452*, 2014. (Cited on page 4.)
- 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 6.)
- Benjamin W. Pugsley, Petra Sedláček, and Vincent Sterk. The nature of firm growth. *Working Paper*, 2017. (Cited on page 35.)
- 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 10.)
- T. Kirk White. Recovering the item-level edit and imputation flags in the 1977-1997 Census of Manufactures. *Census Discussion Paper CES 14-37*, September 2014. (Cited on page 48.)
- T. Kirk White, Jerome P. Reiter, and Amil Petrin. Imputation in U.S. manufacturing data and its implications for productivity dispersion. *Review of Economics and Statistics*, 100(3):502–509, July 2018. (Cited on pages 24, 31, and 47.)

Appendix

A Data and Measurement

A.1 Constructing the Full Census Sample

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

In principle, the Census covers all existing 300-350k establishments in the manufacturing sector. We only consider those establishments that are not administrative records and are in the “tabbed sample,” a distinction Census started in 2002. Non-tabbed establishments are considered by Census to be not really active 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`¹⁹, which stays with the establishment from its birth to its death. That variable is available as a consistent identifier throughout all years.²⁰ 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,

¹⁹All variable abbreviations entirely reflect our own choice.

²⁰Although the `LBDNUM` was created for the Longitudinal Business Database (LBD), which starts in 1976, `LBDNUM`s have been created for the Economic Censuses and the Annual Survey of Manufactures before then as follows: (1) if the plant exists in 1976 or later, Census uses the assigned `LBDNUM` from the later year; (2) if the plant died prior to 1976, Census assigns the same `LBDNUM` (just a made up number that will not conflict with existing `LBDNUM`) to all appearances of that establishment before 1976.

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

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 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 costs, which can broadly be interpreted as benefits. Benefits contain involuntary labor costs (item ILC) 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, this is what Census attempts to do. Yet, it is not certain that all temporary help services or leased employment, were 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 this trend, Census decided in 2002 to sample permanent and leased employment separately. Before then, no specific instructions were given to establishment whether or not to include leased employment in their compensation variables. Given this discrepancy, Census has studied the before/after patterns of employment and labor costs; they concluded that the majority of establishments interpreted the question to include all types of workers and their compensation. If this assessment is correct, our labor share measures would capture all labor costs throughout the sample period. Otherwise, we would be missing a portion of the labor compensation before 2002 and therefore underestimate establishment-level labor shares in that early time period. Given Census’s before/after analysis, this possible missing labor compensation is likely small. Yet, even if one were to assume that it is significant, the underestimation of pre-2002 labor costs would imply that the actual fall in the manufacturing labor share would be even more pronounced than what we currently report and analyze. Hence, we can view our empirical results as a lower bound on the manufacturing labor share decline.

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.

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

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

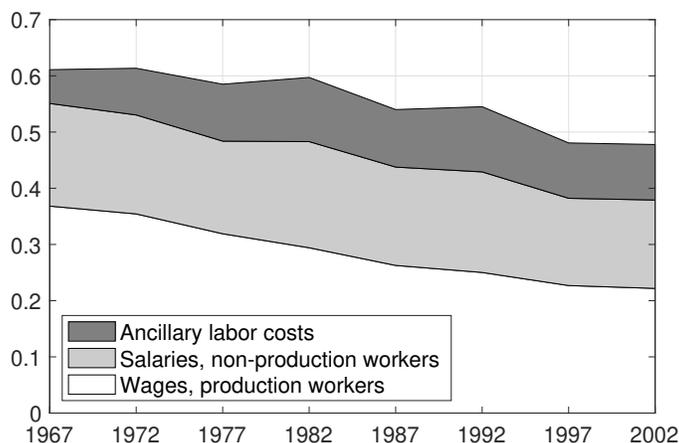
²²Ongoing research in finance is concerned with the rising share of deferred compensation in total labor compensation, see [Eisfeldt, Falato, and Zhang \(2018\)](#).

share components that are shown in Figure A.1 and Table A.1:

$$\lambda_t = \underbrace{\frac{w_t^{pw} L_t^{pw}}{Y_t}}_{\text{Wage bill}} + \underbrace{\frac{w_t^{npw} L_t^{npw}}{Y_t}}_{\text{Salaries}} + \underbrace{\frac{w_t^{ben} L_t^{ben}}{Y_t}}_{\text{Ancill. labor costs}} . \quad (18)$$

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

Figure A.1: 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 wage bill of production workers as a share of value added not declined at all, the manufacturing labor share would have stayed more or less constant (-0.3 ppt per decade).

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

Table A.1: Dynamics of Labor Share Components per Decade

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 (19) applied to the three types of labor compensation listed in Equation (18). 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.

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 adjustment of materials inventory.
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 value added too volatile over time and too dispersed across establishments because it would include a portion of unmeasured fluctuations in intermediate inputs.²³ Since value added plays an important role in the dynamics and aggregation of labor shares, this matters.
- 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 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.

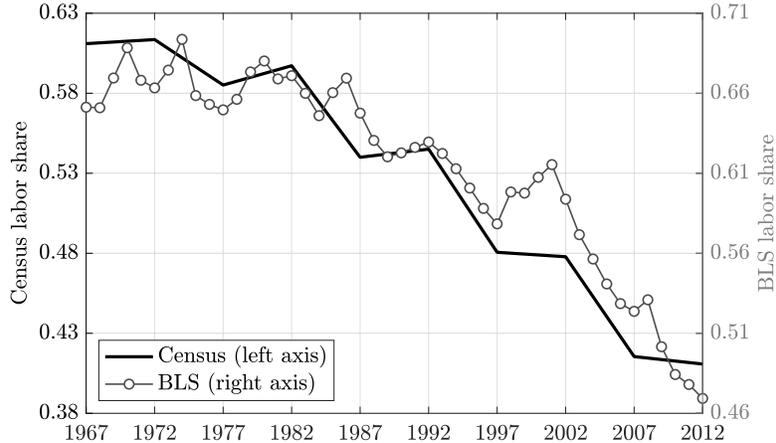
We therefore subtract the industry-year-specific ratio of purchased services to sales from establishment sales.²⁴ This avoids outsourcing contaminating our measure of the labor share and its time series behavior, but it does not impact within-industry reallocation dynamics

²³Though this is not a problem in this paper, treating resales as an intermediate input would cause biases in gross output production function estimates (see Kehrig (2015), p. 41).

²⁴Autor, Dorn, Katz, Patterson, and Van Reenen (forthcoming) pursue a similar strategy.

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 8 ppt, a gap that remains stable over time (see Figure A.2).

Figure A.2: The Labor Share in U.S. Manufacturing



Note: The solid black line (left scale) represents the manufacturing labor share λ_t in the Census panel as calculated in Equation (1); the thin grey line with balls represents the labor share in the manufacturing sector as calculated from BLS data.

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:

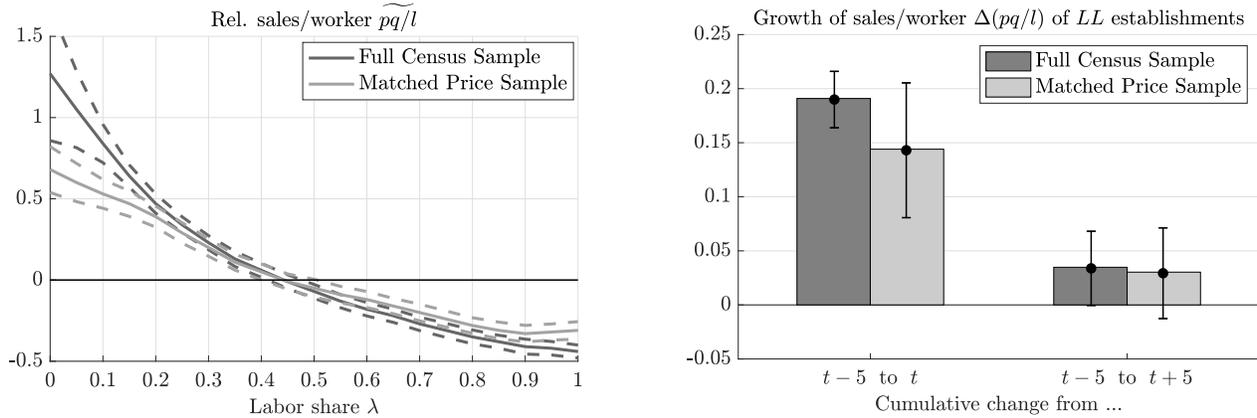
- **R__**: Any observation starting with **R** denotes reported values. Of these, we keep those that were not replaced with an imputed value, in particular:
 - **RC**: analyst correction of reported value,
 - **RG**: goldplated observation (due to analyst information “known” to be of such high quality that any imputation would worsen data quality),
 - **RN**: reported value just corrected for obvious rounding errors;
 - **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 .

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.3](#), we run a non-parametric regression analogous to [Equation \(12\)](#) 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

Figure A.3: 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.

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.3 we display the relative sales-per-worker dynamics of *LL* establishments versus Non-*LL* establishments. The approach is analogous to (16) and (17): 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 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.

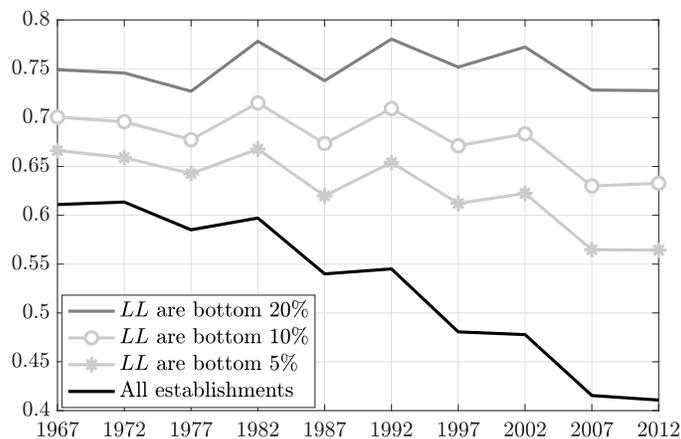
B Defining *LL* Establishments

B.1 Choosing the Cutoff

This section considers the selection criteria for the *LL* establishments that underlie the dynamic analysis in Sections 4 and 5. While any choice of a cutoff is to a certain degree arbitrary, our aim is to isolate the portion of the labor share distribution that plays a central role in the manufacturing

labor share decline. In Figure B.1, we plot the manufacturing labor share once we leave out establishments as the bottom 20%, 10%, 5% and 1% in a given industry and year.

Figure B.1: The Manufacturing Labor Share without Differently Defined *LL* Establishments



As the figure makes it clear, picking a labor share cutoff of 20% in a given 3-digit NAICS industry implies that the remaining 80% of establishments has a stagnant labor share. Any trend estimate is insignificantly different from zero. Choosing a more restrictive cutoff of 10% labor shares, however, shows that the manufacturing labor share of the 90% Non-*LL* establishments declines by 6 ppt. This means that establishments between the 10th and 20th percentile contribute materially to the manufacturing labor share decline. Not counting them as *LL* establishments would lead us to miss some quantitatively relevant dynamics. This is even more evident if only the bottom 5% of establishments are define as *LL*; the manufacturing labor share decline among their complementary set of Non-*LL* establishments would be 10 ppts.

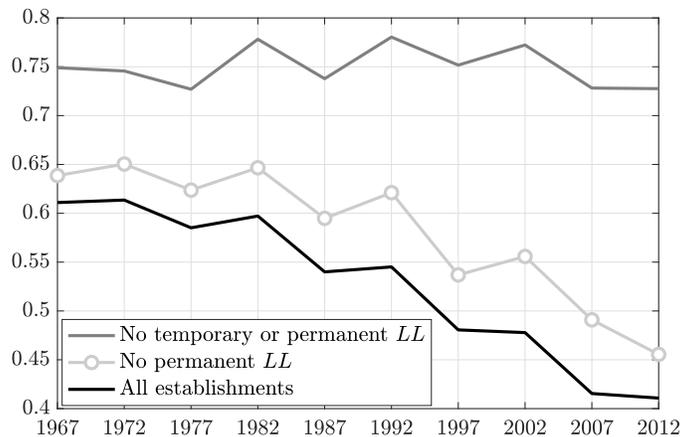
B.2 Permanent versus Transitory *LL* Establishments

In Section 4.6, 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 9, 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 $t - 5$ 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.2. 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.

Figure B.2: The Role of Temporary and Permanent *LL* Establishments



B.3 The Magnitude of Sales Spikes

The V-Shaped pattern documented in Section 4.6 implied that *LL* establishments on average experience a 18 ppt drop in their labor share. In this subsection, we show that the magnitude of the sales increases that is consistent with that impact is sensible.

Consider the growth rate of establishment *i*'s labor share:

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

As we show in Figures 6 and 10 of the paper, *LL* establishments achieve a lower labor share through a stark increase in value added, *PY*, while leaving wages, *W*, and employment, *L*, almost unchanged. A drop in the labor share by 18 ppt (weighted estimates in Figure 9) corresponds to a growth rate of -46.3% which is almost entirely explained by the increase in value added (growth of 45.3%).

While a value added growth rate of 45.3% may sound suspiciously large at first, we show next that it is in fact reasonable. As a first step, consider that value added is defined as:

$$\text{Value Added} = \text{Sales} - \text{Intermediate Inputs}.$$

Note that the level of value added is additive in sales and materials, so their growth rates are not additive. Thus, one has to multiply sales growth and materials growth by their share in value added. In our data, sales of *LL* establishments is 2.57 times as large as value added on average, while materials use is 1.57 times as large. When we repeat the estimation of Equations (16) and (17) in the paper (V-shaped regressions) for the sales and materials growth rates of *LL* establishments relative to their peers, we estimate that their sales to increase by 19.1%, while their intermediate input use increases by only 2.4%:

$$\begin{aligned}
\underbrace{\Delta \log \lambda_{it}^{LL}}_{-46.3\%} &= \underbrace{\Delta \log W_{it}^{LL}}_{-1\%} + \underbrace{\Delta \log L_{it}^{LL}}_{\approx 0\%} - \underbrace{\Delta \log P_{it}^{LL} Y_{it}^{LL}}_{45.3\%} \\
&= \underbrace{\Delta \log W_{it}^{LL}}_{-1\%} + \underbrace{\Delta \log L_{it}^{LL}}_{\approx 0\%} - \underbrace{\Delta \log (Sales_{it}^{LL} - Int_{it}^{LL})}_{45.3\%} \\
&= \underbrace{\Delta \log W_{it}^{LL}}_{-1\%} + \underbrace{\Delta \log L_{it}^{LL}}_{\approx 0\%} - \underbrace{\Delta \log Sales_{it}^{LL} \frac{Sales_{it}^{LL}}{P_{it}^{LL} Y_{it}^{LL}}}_{19.1\% \times 2.57} + \underbrace{\Delta \log Int_{it}^{LL} \frac{Int_{it}^{LL}}{P_{it}^{LL} Y_{it}^{LL}}}_{2.4\% \times 1.57}
\end{aligned}$$

This means that, in order to explain the admittedly strong growth in value added of 46.3%, sales only have to increase by 19.1%. Moreover, it is important to remember that these values correspond to growth rates over five-year periods; the annualized growth rate is only about 3.5%. We see it as plausible that *LL* establishments grow their sales 3.5% faster per year than their Non-*LL* peers to achieve their low labor share over a five-year span.

C Establishments vs. Firms

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 establishments likely sort into the firms along labor share. *LL* establishments, in particular, would sort into the same firms, which would then become *LL* 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 9. If *LL* establishments are evenly distributed across firms, however, we would expect firm-level labor shares to be much more smooth and as establishment-level labor share dynamics get diversified away by the firm.

To that end, we aggregate labor cost and value added across a *LL* establishments within the same firm (defined by FIRMID) to compute firm-level labor shares. In principle, FIRMIDs stay with the same firms, but we follow [Dent, Pugsley, and Wheeler \(2018\)](#) to correct false identifier changes (say, after a firm changes legal form or transitions from single-unit to multi-unit firm) and indicate changes where there should be one (say, after mergers and acquisitions). This correction will only matter for the dynamic analysis, where we have to follow firms over time. Following other research, we break firms along their 3-digit NAICS industry codes to avoid subsectoral differences drive labor shares of extremely diversified firms.²⁵

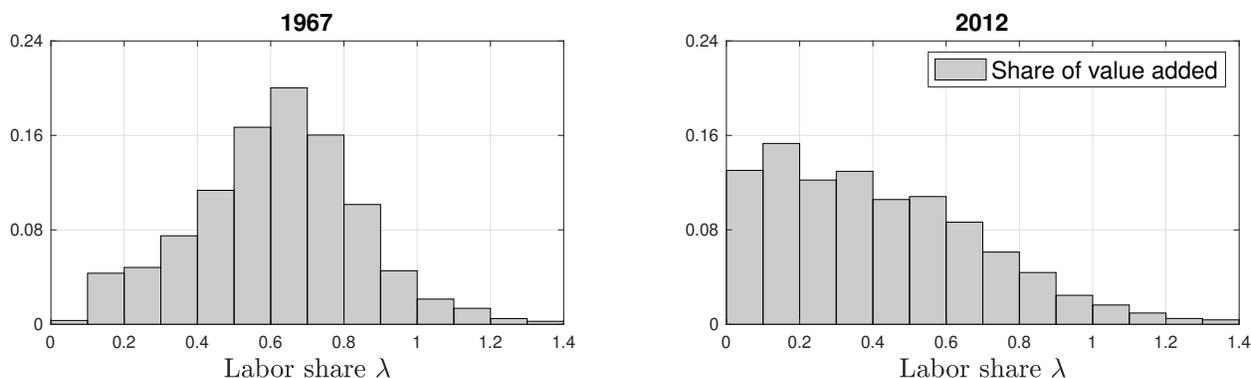
C.1 The Reallocation of Value Added across Firms

In what follows, we repeat the main empirical exercises from Section 4 but at the firm instead of establishment level. Comparisons between these two levels of aggregation can teach us about the nature of labor share dynamics. For example, if we fail to detect significant reallocation of economic

²⁵Using a more restrictive definition of a firm only within the same 4-digit or even 6-digit NAICS industry code would deliver similar results (NAICS-4) or stronger (NAICS-6) results.

activity towards low-labor share firms, then we would conclude that reallocation is mostly a within-firm phenomenon.

Figure C.1: The Reallocation of Value Added between Firms



Note: This figure depicts the firm-level evidence analogously to Figures 1 and 3 (right column).

The firm-level reallocation displayed in Figure C.1 shows an overall pattern that is similar to the establishment-level reallocation that was depicted in 3: in 1967, more than half of value added is being produced by firms with a labor share between 50 and 80 ppts, while there is little output accounted for by firms with a very low labor share. By 2012, in contrast, most of value added has been reallocated to low-labor share firms. Specifically, half of manufacturing value added is being produced by firms with a labor share of less than 37 ppts; the analog number for establishments was 32 ppts. So even if the reallocation is not as dramatic as it is for establishments, it remains very strong. We conclude that most of the reallocation of value added takes place between rather than within firms.

C.2 The Joint Dynamics of Labor Share and Value Added

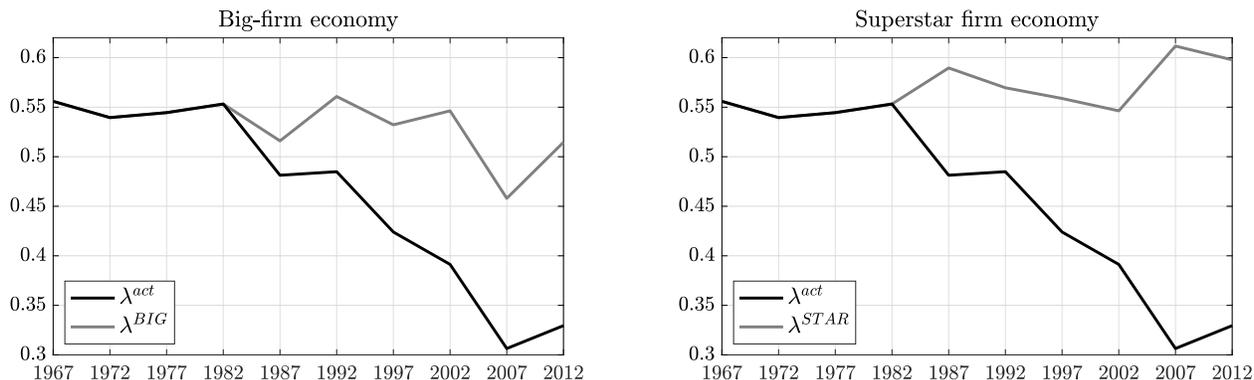
Next, we analyze the role of compositional changes in driving the firm-level reallocation, as we did in 4.3.2 for establishments. While we found that neither differential labor share dynamics by size (the big-player scenario) nor differential growth by labor share (the superstar scenario) played a role at the establishment level, this may yet be true at the firm level. For example, superstar firms may very well drive the manufacturing labor share decline even if only a subset of their establishments are *LL* establishments, as long as this set rotates within the firm.

Similar to our exercises in Section 4.3.2, we construct counterfactual manufacturing labor share measures λ_t^{BIG} and λ_t^{STAR} based on firm-level labor and market shares. If the dynamics of these counterfactual labor shares, which are plotted in Figure C.2, were identical to the actual labor share, then the hypothesis that large firms or superstar firms drive the manufacturing labor share decline might have empirical support. In that case, compositional changes at the firm level would play a role even though they do not across establishments.

The counterfactual labor shares in Figure C.2 look quite similar to their establishment-level counterparts in Figure 5: there is no contribution from firms with an initially low labor share (the initial superstars) and only a very limited contribution of initially large firms. Firms that were already large in 1982²⁶ lowered their labor share enough to make the manufacturing labor share

²⁶Again, we confirm that taking the average of 1977, 1982 and 1987 as the “initial level” of labor share or market shares does not make a difference.

Figure C.2: The Limited Role of Big Firms or Superstar Firms



Note: This figure depicts the firm-level evidence analogously to Figure 5.

decline by 7 ppts, which is about 3 ppts more than the contributions of initially large establishments, but still far from the 21 ppt-decline in the actual manufacturing labor share. For the superstar counterfactual, we again observe an increase in the counterfactual labor share, suggesting that firms with initially low labor shares did not experience strong enough growth to contribute to the manufacturing labor share decline on their own.

C.3 The V-Shape Labor Share Dynamics of *LL* Firms

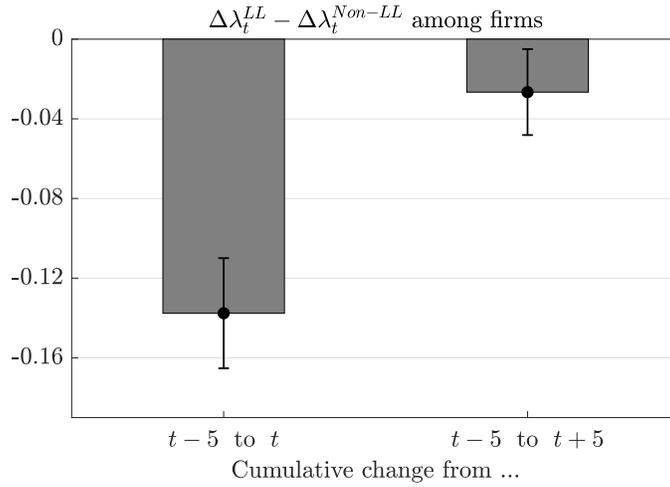
Perhaps one of the most important – and most surprising – findings of Section 4 was the V-shaped labor share pattern of *LL* establishments. Labor shares are typically low for a period of five to eight years, and about almost 60% of *LL* establishments are no longer *LL* establishments five years later. In theory, it is not obvious that the same pattern should hold at the level of the firm. For example, imagine a firm whose establishments alternate in *LL* status: half of them experience in period t temporary demand shocks specific to the products they produce, while for the other half this occurs at $t + 5$. Alternatively, this rotating pattern may occur if a vertically or horizontally integrated firm shifts profits from one establishment to another using transfer prices, possibly for tax purposes. In both cases, the V-shaped labor share pattern of establishments would wash out at the firm level. On the other hand, demand shocks may instead stem from a shift in the firm’s marketing strategy or brand appeal. In that case, labor shares across establishments within the same firm would be positively correlated, and *LL* firms, like their establishments, would exhibit a V-shaped labor share pattern.

In order to assess these within-firm dynamics, we define “*LL* firms” analogously to *LL* establishments: their labor share is in the lowest quintile of the firm’s industry in a given year. We then repeat the analysis of (16) and (17) for these *LL* firms and show them in Figure C.3.

Clearly, the V-shaped pattern is still present for *LL* firms even though its magnitude (14.5 ppts) is slightly smaller than for *LL* establishments (18 ppts). The rebound is also similar: between $t - 5$ and $t + 5$, the labor share of a time- t *LL* firm falls by a mere 3.1 ppts relative to that of its peers (2.6 ppts for *LL* establishments). For the unweighted estimates (not disclosed), the V-shapes of *LL* establishments and *LL* firms look equally large. This leads us to two conclusions: First, within-firm reallocation dynamics are not the main cause of the V-shape documented above.²⁷ Second,

²⁷An implication is that transfer prices are not the main driver of the price dynamics documented in Section 4.5.

Figure C.3: Labor Share Dynamics of *LL* Firms



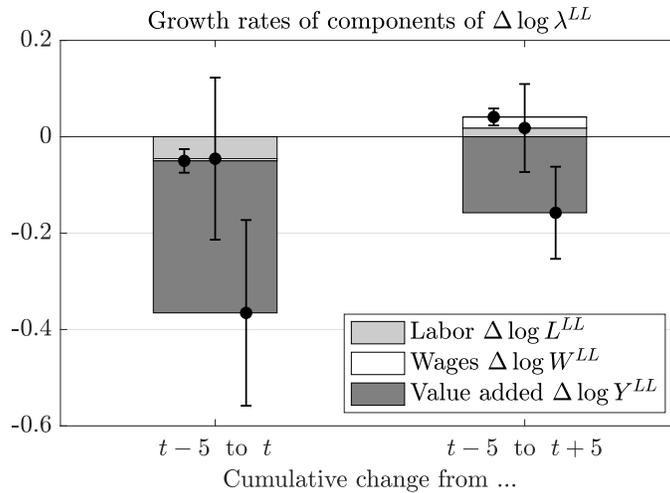
Note: This figure depicts the firm-level evidence analogously to Figure 9.

LL status across tends to co-move positively across establishments within the same firm.

C.4 The Dominant Role of Value Added

Next, we study the relative dynamics of the components of the labor share of *LL* firms, analogously to the exercise of Section 4.6.1. Figure C.4 shows that value added dominates the dynamics of firm-level labor shares, in line with the establishment-level evidence of Figure 10.

Figure C.4: Wage, Employment and Value Added Dynamics of *LL* Firms



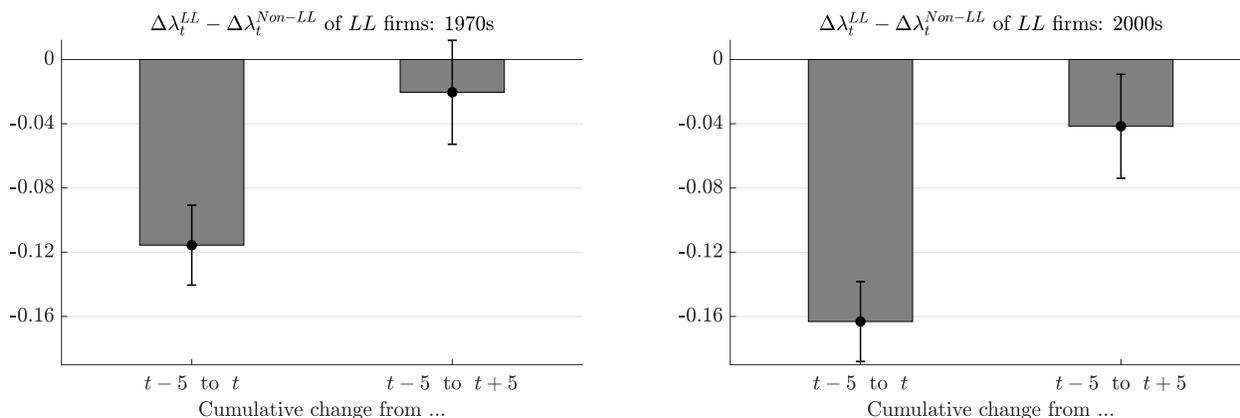
Note: This figure depicts the firm-level evidence analogously to Figure 10.

C.5 Did Firm-Level Shocks Become More Volatile over Time?

Lastly, we repeat the analysis of Section 5 on firm data.

The deepening V-Shape of *LL* firms Figure C.5 shows that similar to the evidence for establishments, firm-level V-shapes have also become deeper over time. While the typical *LL* firm dropped its labor share by 11 ppts in the 1970s relative to its peers (14 ppts for *LL* establishments), this became a 16.6 ppts drop in the 2000s (21 ppts for *LL* establishments). Again, while the firm dynamics are slightly more muted, they are qualitatively and quantitatively very comparable to those of establishments.

Figure C.5: The Deepening of the V-Shape of *LL* Firms



Note: This figure depicts the firm-level evidence analogously to Figure 12.

The increasing disconnect between value added and employment Looking at the evolution of the labor share component dynamics, we can see in Figure C.6 that, in the 1970s, hiring helped buffer the negative contribution of value added to the labor share of *LL* firms and used to significantly contribute to the subsequent rebound. In the 2000s, in contrast, relative employment would on average fall for *LL* firms, though the point estimate is not statistically significant.

In sum, the evidence in this appendix shows that the findings documented in the main body of the paper are not specific to establishments but instead carry through when the analysis is performed with firm data. This indicates that the forces, factors and shocks underlying micro-level labor share dynamics are most likely taking place at the level of the firm.

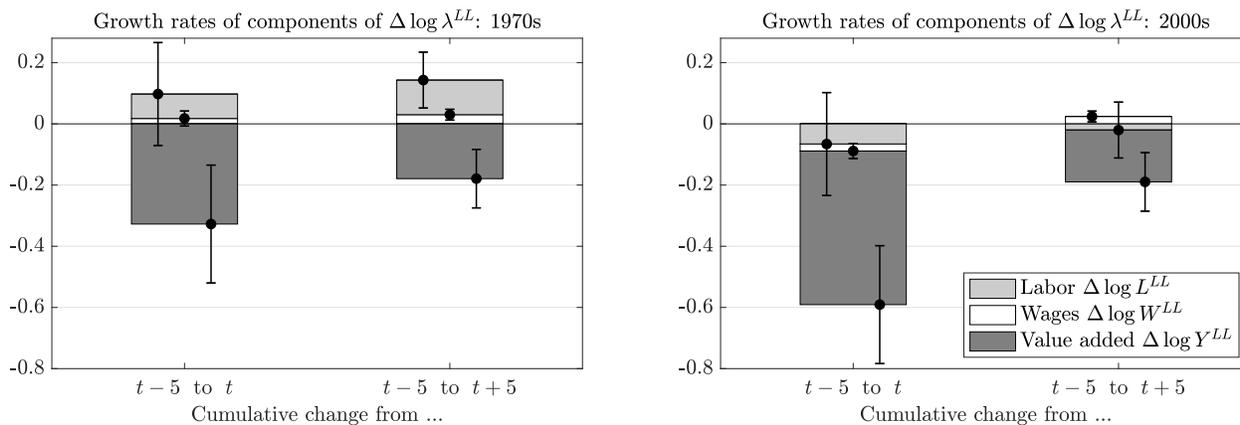
D Robustness on the Importance of Demand

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

D.1 Corroborating Evidence: The Importance of Advertising for Demand

In Sections 4.2-4.5, we have used our conceptual framework to identify demand factors as the only class of shocks that is consistent with all the empirical micro-level facts we documented. While that is helpful in guiding the profession in developing models that incorporate these insights, “demand shocks” remain a fairly abstract notion. Clearly, superstar products such as the Apple

Figure C.6: Wage, Employment and Value Added Dynamics of *LL* Firms



Note: This figure depicts the firm-level evidence analogously to Figure 13.

iPhone or AbbVie’s Humira, the no. 1 blockbuster drug every year since 2015, generate sales because customers value that product rather than efficiency advances over generic competition. But what do businesses actually *do* that makes their demand curve inelastic so they can charge higher prices? The many factors that may engineer a “customer preference shock” to lift demand for an establishment’s products and makes customers more willing to pay a higher price are known to a business. It will likely invest in its brand value, develop and expand its customer base or spend a lot on market research when developing new products. To our knowledge, information about on the relationship between these activities and labor shares is missing in the literature. As it turns out, Census started to collect some data on advertising expenditures since 1997. While this likely represented only one of many types of spending that may shape and drive demand factors, it may still hint at how *LL* establishments manage to convince customers to pay higher prices for their products relative to the competition. In fact, advertising has steadily increased over time and found to matter for firm profits, sales and markups (Gourio and Rudanko (2014)) in general. Next, we show that advertising intensity is significantly related to labor shares, hinting at a potential role in driving the micro-level dynamics we documented earlier.

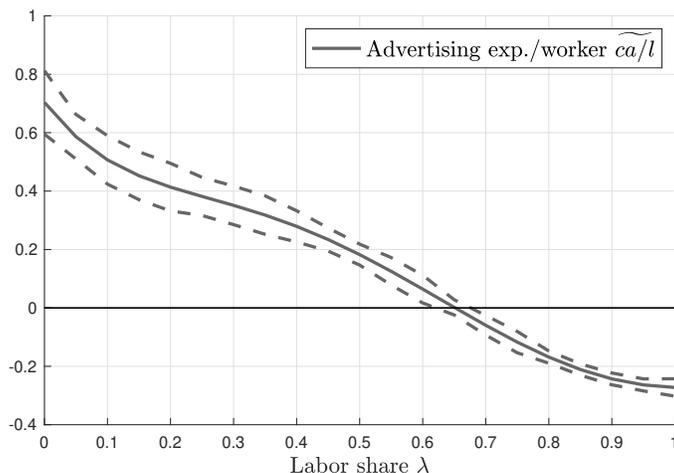
First, it should be mentioned that the dataset only contains advertising expenditures at the establishment level, while such expenditures at the headquarter or firm level are missing along with other aspects of customer-related marketing investments.

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 (11). 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 D.1.

The figure reveals a clear pattern: low-labor share establishments spend significantly 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 more on advertising per employee than the average plant in its sector and region, while high-labor share establishments ($\lambda = 1$) spend about 25% less.

While it should not be interpreted as causal at this point, we view the relationship between ad-

Figure D.1: Cost of Advertising per Employee and Labor Shares



Note: The figure displays the cross-sectional differences in relative cost of advertising expenditures per employee $\widetilde{ca/l}$ against the labor share. All relative measures denote log-point differences vis-à-vis their peers as defined in Equation (11). Dashed lines denote 95% error bands.

vertising expenditures and labor shares as corroborating evidence about the importance of demand shocks.

D.2 The Unimportance of Industry, Regional, Legal Factors

To test for industry and/or geographical composition effects, we decompose the manufacturing labor share decline into within- and between-groups components using Equation (19):

$$\Delta\lambda_t = \underbrace{\sum_j \Delta\lambda_{jt}\omega_{jt-5}}_{\text{Within adjustment}} + \underbrace{\sum_j \lambda_{jt-5}d\omega_{jt}}_{\text{Between reallocation}} + \underbrace{\sum_j \Delta\lambda_{jt}d\omega_{jt}}_{\text{Residual}} \quad (19)$$

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 D.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 per decade 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 output 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 D.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

Table D.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
<i>A. NAICS-3 industries</i>			
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
<i>B. Census regional divisions</i>			
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
<i>C. Legal form of organization</i>			
Within-LFO adjustment	-6.3	+1.1	-6.6
Between-LFO reallocation	+0.3	-0.6	+0.4
Residual	+0.4	+1.8	+0.0
<i>D. Public vs. private firms</i>			
Within-group adjustment	-5.1	-0.5	-7.9
Between-group reallocation	+0.2	-0.5	+0.5
Residual	+0.1	+0.0	+0.1

Note: Results from the shift-share decompositions as defined in (19) applied to industries (Panel 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.

when adding the residual term. An analogous analysis at the state level shows similar results.²⁸

Next, we study the effect of the legal form of organization. The 1980s saw the emergence of new legal forms of organization such as S-corporations. This has been studied in many papers, (see, among others, [Dyrda and Pugsley \(2019\)](#)), and it is plausible to think that the 1980s tax reform may have had an effect on the labor share if pass-through entities have diverted some labor incomes of proprietors into profits passed through to firm owners. We combine our baseline dataset with the Standard Statistical Establishment List (SSEL), which contains the legal form of organization for single- and multi-unit firms, including the break-down of corporations into C- and S-corporations.²⁹ While S-corporations do become much more important, Panel C. suggests they did not play a role in the decline of the manufacturing labor share. Most of the decline of the labor share occurred within the same category of legal forms of organization rather than resulting from a shift from sole proprietorships and C-corporations to S-corporations.

Lastly, we study if a shift of economic activity to publicly traded firms matter for the labor share decline. Those firms likely face less financial frictions and can more easily build capital. With

²⁸Estimating if establishments are more likely to become a superstar once the state enacts right-to-work legislation, we find a statistically significant but economically small effect.

²⁹We thank Benjamin Pugsley for helpful discussions about measuring the legal form of organization of various types of firms.

higher capital intensity, they may have naturally a lower labor share. Again, we find that the labor share decline occurs among both publicly traded and privately held firms (see Panel D.).