# Labor Market Power and Technological Change in US Manufacturing

Ivan Kirov
James Traina*

October 2023

**Abstract**

We estimate time-varying plant-level production functions with Census microdata to separately identify labor and product market power in the US manufacturing sector. Wage markdowns rose substantially from 1 in 1972 to 2 in 2014, while price markups stayed flat at 1. Wage markdowns rose because marginal-revenue-product growth speeds up, not because wage growth stagnates. In local labor markets, wage-markdown growth is uncorrelated with employer-concentration growth. We document strong associations with direct measures of information and communication technologies and indirect measures of management and automation technologies. Altogether, the evidence points to technological threat as a key driver of labor market power.

JEL Codes: D24, D43, J42, L11, O33
Keywords: Labor Market Power, Markdowns, Markups, Technological Change, Labor Share

*Kirov: Analysis Group, ikirov@uchicago.edu. Traina: Federal Reserve Bank of San Francisco, traina@uchicago.edu. We are grateful for the guidance of our advisors Erik Hurst, Chad Syverson, Ali Hortaçsu, Brent Neiman, and Luigi Zingales. We thank Brian Albrecht, Carter Davis, Zach Flynn, Amit Gandhi, Luis Garicano, Mike Gibbs, John Grigsby, Arshia Hashemi, Shohini Kundu, Ben Marrow, Jordan Rosenthal-Kay, Alex Sollaci, and Uyen Tran, and participants in assorted seminars and workshops for comments. We appreciate Chang Hsieh, Huiyu Li, Pete Klenow, Pablo Ottonello, and Tom Winberry for data support, and the Chicago Booth Stigler Center for financial support. All expressed views are those of the authors and not of the US Census Bureau, Analysis Group, or the Federal Reserve System. Results have been reviewed to ensure no confidential information is disclosed.

The US labor share of income fell from 64% in the mid-1980s to 58% in 2012 (Elsby et al., 2013; Karabarbounis and Neiman, 2014). This fall was sharper in manufacturing, where the labor share fell from 57% to 41% (Kehrig and Vincent, 2021). Over the same period, average productivity decoupled from average pay (Bivens and Mishel, 2015; Stansbury and Summers, 2017), with the biggest gap in technology-intensive manufacturing (Brill et al., 2017).

We explore an emerging hypothesis for these trends: rising labor market power.[1] We apply contemporary methods in production function estimation to US manufacturing microdata to measure labor market power over production workers. Manufacturing holds particular interest for two reasons: (1) The sector is large, still accounting for 10% of US employment even after its sharp fall in labor share, and (2) administrative data are available for comparable outputs and inputs from representative samples spanning over four decades.

We define labor market power $\Delta$ as the ratio of a firm's marginal revenue product of labor $R_L$ and its wage $W$: $\Delta = \frac{R_L}{W}$. In an undistorted competitive benchmark, this wage markdown equals 1. With labor market power distortions, firms restrict their labor-input choice, creating a wedge on the margin between revenue productivity and pay. Firms with product market power restrict intermediates and labor symmetrically, but firms with labor market power restrict labor differentially. If intermediate-input markets are undistorted, and labor-input markets are distorted by market power only, we can write $\Delta$ as the wedge between a cost-minimizing firm's intermediates and labor choices:

$$\Delta = \frac{R_L}{W} = \frac{f_l}{f_m}\frac{CM}{WL} \tag{1}$$

where $\frac{f_l}{f_m}$ is the elasticity ratio for labor to intermediates, and $\frac{CM}{WL}$ is the cost ratio for intermediates to labor. Though recovering the elasticity ratio $\frac{f_l}{f_m}$ requires production function estimation, the cost ratio $\frac{CM}{WL}$ is measurable with public data.

[1] Labor market power is often called monopsony power after Robinson (1933). We use the more general term because our framework nests other forms of competition, e.g., oligopsony.

The solid black line in Figure 1 plots the growth in the sector's $\frac{CM}{WL}$ cost ratio using the NBER-CES Manufacturing Industry Database. We normalize the series to 1 at the start of the sample. The ratio grows from 1 in 1958 to 1.5 in 1980, climbs to 1.7 over the 1980s and 1990s, and surges to 2.7 by 2014. For comparison, the dashed gray line plots the remarkably stable ratio of revenue to intermediate-input costs $\frac{R}{CM}$, the cost ratio for product market power—the same cost-minimization problem yields the equation for price markups: $\mathcal{M} = f_m \frac{R}{CM}$. Under the benchmark of equation 1, the extraordinary growth in the $\frac{CM}{WL}$ cost ratio suggests a significant rise in labor market power. This paper asks and answers the key resulting question: Do trends in the output elasticities $f_l$ and $f_m$ offset it?

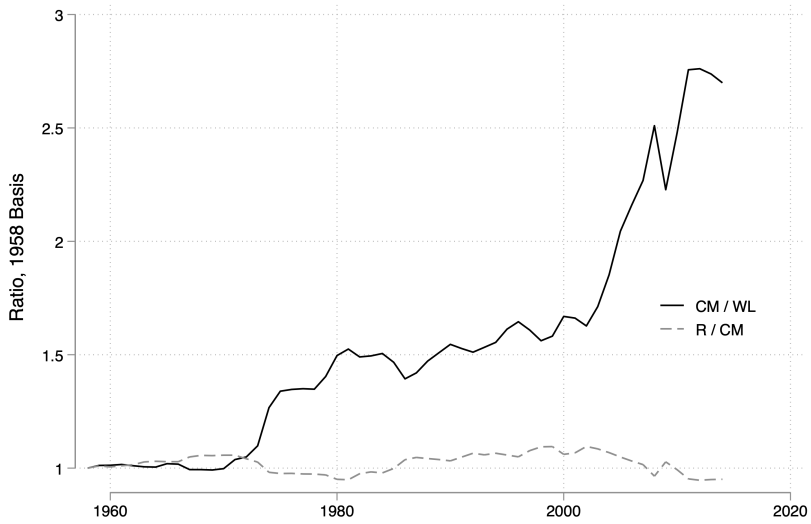


Figure 1: Cost-Ratio Trends Suggest Rising Labor Market Power

Authors' calculations using the NBER-CES Manufacturing Industry Database, March 2021 file. Intermediate-input cost $CM$ is the total expenditure on energy and materials; labor-input cost $WL$ is the total expenditure on production workers; and revenue $R$ is the total value of shipments. Labor costs exclude nonproduction workers, such as technology professionals and managers. The data contain annual industry-level measures of outputs and inputs from 1958 to 2018, derived from about 300,000 plants surveyed in the Census of Manufactures for years ending in 2 and 7, and 50,000 plants surveyed in the Annual Survey of Manufactures for all other years.

The answer: They do not. Using the underlying microdata of Figure 1, we estimate manufacturing production functions to recover $\Delta$ at the plant-year level. We focus on measuring labor market power over production workers—those directly engaged in manufacturing tasks such as fabricating, processing, assembling, and inspecting. For each narrowly defined 6-digit NAICS industry, we allow for heterogeneity in the elasticity ratio with a separate four-input translog specification[2] that lets both $f_l$ and $f_m$ vary by plant and over time according to input intensity, up to weak shape restrictions. Our novel econometric approach addresses the nonidentification challenges of Gandhi et al. (2020) and Bond et al. (2021) using the methods of Flynn et al. (2019) and Kirov et al. (2023). Relative to other approaches (e.g., Azar et al. (2022) or Lamadon et al. (2022)), we impose minimal structure on the nature of labor market competition, crucial if we want to understand how conduct evolves.

We find labor market power in the US manufacturing sector is currently high and has soared since the 1970s: firms paid production workers their marginal revenue product in 1972, but only half this amount by 2014. By contrast, the average price markup stays flat at 1. The average labor wedge was approximately 1 until 1990, implying labor markets were competitive. It then rose to 1.2 over the 1990s (implying $\frac{1.2-1}{1.2} = 17\%$ lower wages). The year 2002 marks an inflection point, from which the wedge widens to 2 by the end of our sample in 2014. Trends in the elasticity ratio $\frac{f_l}{f_m}$ do not offset trends in the cost ratio $\frac{CM}{WL}$; in fact, the elasticity ratio is stable at its time-series average of 0.18 with a standard deviation of 0.01.

Are these findings driven by outsourcing or offshoring? Household-survey and imported-commodities data show potential mismeasurement of the $\frac{CM}{WL}$ cost ratio from these factors is not large enough to offset our estimated rise in labor market power. More generally, alternative explanations must differentially generate trends in labor relative to intermediates, with inflection points at the same time and in the same direction. Outsourced service jobs disappear from

[2]The four inputs are intermediates (energy and materials), production workers, nonproduction workers, and capital (equipment and structures); the translog specification is a second-order approximation to any production function.

plant-level labor bills, but these workers still report their industry as manufacturing in household surveys. The share of manufacturing workers in service jobs increases from 11% in 1990 to 14% in 2010, but conservatively adding these workers back into the cost ratio's denominator only meaningfully lowers the level of markdowns, not the trend. We analyze industry import data similarly, but assuming a double effect of undermeasured labor and overmeasured intermediates. And we find a similar answer: moving all same-industry import expenditures from $CM$ to $WL$ does not meaningfully alter the trend, consistent with the relative importance of other factors over offshoring shown in other settings (Goos et al., 2014).

When decomposed into productivity $R_L$ and pay $W$, the overall increase in $\Delta$ comes from an acceleration of $R_L$ starting in the 1990s. This process speeds up further around 2002, coinciding with the boom in information and automation technologies. Our result is at odds with the conventional wisdom of stagnant wages and suggests a large role for technological change and a rising marginal cost curve for labor. Rising marginal cost curves complement the literature's current focus on falling labor demand (Autor et al., 2013; Autor and Salomons, 2018; Fort et al., 2018; Charles et al., 2019; Acemoglu and Restrepo, 2019). Manning (2006) emphasizes the firm's effective marginal cost of labor can slope up with diseconomies of scale in hiring. In the model, the firm's labor supply curve depends on both wages and recruitment intensity; labor market power arises when the cost of recruiting an extra worker is increasing in total employment, such as when new technologies demand new types of workers (Blatter et al., 2012). This skill mismatch is consistent with the fact that manufacturing saw a differentially large increase in job openings at the same time as a differentially large decrease in employment, as often reported in the popular press (Elejalde-Ruiz, 2016; Sussman, 2016).

We exploit our microdata plant-year structural estimates to confirm labor wedges negatively predict labor shares, and can therefore help explain the collapse of the manufacturing labor share. Panel regressions show a 10% increase in $\Delta$ is associated with a 1-percentage-point decrease in a plant's labor share,

working through lower employment. To get a sense of size, applying this estimate to the $\Delta$ time series suggests the doubling of wage markdowns reduced the manufacturing labor share by 10 percentage points, half the total decline. Although this calculation abstracts from general equilibrium effects, it indicates the rise in labor market power is an important contributor to the fall of the labor share. Turning to the role of spatial competition, when we regress local labor market power on local labor market concentration, we find a positive statistical relationship in levels but not in changes. This finding supports existing work showing labor market concentration reduces wages (Rinz, 2022; Berger et al., 2022; Hershbein et al., 2019; Azar et al., 2022; Benmelech et al., 2022), but calls for skepticism in inferring secular trends in labor market power from cross-sectional concentration variation. In fact, an analysis of variance reveals a plant's industry explains seven times more variation in markdowns than a plant's commuting zone.

We find technological change is a significant source of labor market power. In select years, the US Census Bureau surveys plants on their computer and communications equipment expenditures, which we exploit to test the role of technological change in our microdata estimates. The elasticity of the labor wedge with respect to new computer and communications expenditures per worker is about 0.1 each, so that a 10% increase in technology intensity predicts a 1% increase in wage markdowns. Higher wage markdowns are also associated with more managerial employees and deeper capital stocks, proxies for more advanced technologies (Atalay et al., 2014).

How might technological change lead to labor market power? One possibility is skill mismatch viewed through the lens of Manning (2006). An alternative possibility is technological threat—when firms gain the option to adopt new technologies that can substitute for workers in production, they also gain a stronger bargaining position. Leduc and Liu (2023) applies this reasoning to business cycles, finding the threat of automation dampens wage adjustments and amplifies unemployment fluctuations. One might also conceptualize, as labor gets displaced by new technologies, firms also accrue threats of replacing

a remaining worker with someone who just got laid off.

In the last part of our paper, we explore the relationship between labor market power, unionization, and offshorability using 1990 to 2010 long-difference regressions. At the subsector level, a 10-percentage-point decrease in unionization rates predicts a 17% increase in markdowns; the corresponding relationship with offshorability rates is statistically indistinguishable from zero. Paired with the fact that US manufacturing unionization rates have steadily declined since the 1970s, our evidence points to unions as a mediator of labor market power.

Our paper links the literatures on labor markets, market power, and technological change. In fact, our results underscore technological change as a key driver of labor market power over workers. Our main finding supports the broader concern of labor market power in the US economy (Shapiro, 2019; Stansbury and Summers, 2020). This concern stands somewhat in contrast to that of product market power, for which evidence is mixed;[3] Our research shows the average markup in US manufacturing is 1 for four decades. Autor and Dorn (2013) and Acemoglu and Restrepo (2019) explore how technological change polarizes labor markets and displaces routine workers, and Autor et al. (2020) connect related changes to product market concentration. We argue technological change also alters the structure of labor markets, further affecting worker outcomes by increasing labor market power.

# 1 Measuring Markdowns

This section outlines a model of labor market power along the lines of Dobbelaere and Mairesse (2013). Our starting point is to assume firms can flexibly adjust intermediate and labor inputs. If the intermediate-input markets are competitive, cost minimization implies when a firm undersupplies intermediates relative to the competitive benchmark, we can infer a market power

[3]See De Loecker et al. (2020), Traina (2018), and Hall (2018) for a wide range of markup trend estimates; or see the reviews in the Journal of Economic Perspectives, Summer 2019 Markups Symposium (Basu, 2019; Syverson, 2019; Berry et al., 2019).

wedge and a resulting markup in the product market. Pricing power in the labor-input market generates an extra wedge above and beyond the markup. Following Dobbelaere and Mairesse (2013), we define labor market power as this extra wedge, equivalent to the ratio of the wedge implied by intermediates (which only has a markup) and the wedge implied by labor (which has both a markup and a markdown).

A few notes on notation are in order. Throughout the paper, we use uppercase letters to denote levels and lowercase letters to denote logs, so that $v = \log V$ for variable $V$. Derivatives are in subscript form, so that $F_V = \frac{\partial F}{\partial V}$ for function $F$. Combined, these conventions imply a simple form for elasticities: $f_v = \frac{\partial f}{\partial v} = \frac{\partial \log F}{\partial \log V} = \frac{V \partial F}{F \partial V}$. Finally, because we apply our conceptual framework to annual plant data, all variables are implicitly understood to be at this level.

## 1.1 Labor Wedges from Cost Minimization

This section develops a deterministic market-power framework following Dobbelaere and Mairesse (2013), De Loecker and Warzynski (2012), and Morlacco (2020), who build on Hall (1988). Consider a firm $i$ at time $t$ that uses intermediates $M_{it}$, production labor $L_{it}$, nonproduction labor $N_{it}$, and capital $K_{it}$ to produce output $Q_{it}$. The firm faces competitive intermediate-input markets, so it can buy as much as it likes at price $C_t$. The firm potentially has market power in product and production labor markets. For the latter, choosing lower labor inputs also means paying a lower equilibrium wage. Nonproduction labor and capital are predetermined, so the short-run cost-minimization problem is over $M_{it}$ and $L_{it}$ only.

The firm's production function $F$ is given by

$$Q_{it} = A_{it} F(M_{it}, L_{it}; N_{it}, K_{it}) \tag{2}$$

where $A_{it}$ is Hicks-neutral total factor productivity. The only restrictions we need on $F$ is that it's continuous and twice differentiable.

We formulate the firm's short-run cost-minimization problem as:

$$\begin{array}{cc} \min\limits_{M_{it},L_{it}} & C_t M_{it} + W(L_{it})L_{it} \\ s.t. & Q(M_{it}, L_{it}) = \bar{Q}_{it} \end{array} \tag{3}$$

where imperfect competition in the labor market implies wages are a function of labor-input choice, and $Q(M_{it}, L_{it}) = F(M_{it}, L_{it}; N_{it}, K_{it})$ is the short-run output function.

The first-order conditions of the Lagrangian are:

$$\begin{array}{cccc} [M] & C_t & = & \Lambda_{it} Q_{M\,it} \\ [L] & W_{it} + W_{L\,it} L_{it} & = & \Lambda_{it} Q_{L\,it} \end{array}$$

where $\Lambda_{it}$ is the shadow price of output (marginal cost) and $W_{L\,it} L_{it}$ measures labor market power. Now define the price markup as $\mathcal{M}_{it} = \frac{P_{it}}{\Lambda_{it}}$, and the wage markdown as $\Delta_{it} = \frac{R_{L\,it}}{W_{it}}$, where $R_{L\,it}$ is the marginal revenue product of labor for revenue $R_{it} = P(Q_{it})Q_{it}$. Equivalently, we can characterize markups by the firm's inverse residual demand elasticity $p_{q_{it}}$: $\mathcal{M}_{it} = \frac{1}{1+p_{q_{it}}}$; and markdowns by the firm's inverse residual supply elasticity $w_{l\,it}$: $\Delta_{it} = 1 + w_{l\,it}$ (Baker and Bresnahan, 1988; Robinson, 1933). If labor markets are competitive, $w_{l\,it} = 0$, which implies $\Delta_{it} = 1$.

Our wedges are reasonably agnostic to the nature of product demand, labor supply, and competition in either market. Critically, the residual elasticities are not the same elasticities derived from a household's problem—they also depend on the underlying nature of competition, which we leave general. For example, in contestable product markets, the threat of entry makes the incumbents' residual demand curve more elastic than the households' demand curve (Baumol, 1982).

Taken together, our wedges are:

$$\begin{array}{ccccc} \mathcal{M}_{it} & = & \frac{P_{it}}{\Lambda_{it}} & = & \frac{1}{1+p_{q_{it}}} \\ \Delta_{it} & = & \frac{R_{L\,it}}{W_{it}} & = & 1 + w_{l\,it} \end{array}$$

where, respectively, $P_{it}$ and $W_{it}$ are output- and labor-input prices, $\Lambda_{it}$ and $R_{Lit}$ are marginal cost and the marginal revenue productivity of labor, and $p_{q_{it}}$ and $w_{lit}$ are residual elasticity terms.

Combining our wedge definitions with the cost minimization first-order conditions, we have:

$$\begin{array}{cccc} [M] & \mathcal{M}_{it} & = & f_{mit}\frac{R_{it}}{C_t M_{it}} \\ [L] & \mathcal{M}_{it}\Delta_{it} & = & f_{lit}\frac{R_{it}}{W_{it}L_{it}} \end{array} \tag{4}$$

and their quotient:

$$\Delta_{it} = \frac{f_{lit}}{f_{mit}}\frac{C_t M_{it}}{W_{it}L_{it}} \tag{5}$$

In summary, firms buy intermediates in a competitive market. If product markets are also competitive, firms equate average and marginal products of intermediates. If a firm uses fewer intermediates than the competitive benchmark, we infer a product wedge. Such a wedge reduces utilization of all inputs symmetrically. Labor wedges, however, reduce labor utilization only. Comparing the ratio of average to marginal products of labor and intermediates allows us to recover these labor wedges.

## 1.2 Interpreting Wedges as Market Power

Equations 4 and 5 give useful intuition on the nature of labor market power in our model. As a starting point, in the absence of any distortions so that $\mathcal{M} = \Delta = 1$, equation 4 implies the firm chooses its inputs to equate their cost share of revenue with their respective output elasticities. This implication is behind many models in the productivity literature that assume undistorted output and input markets to identify output elasticities and hence production functions and productivity (Syverson, 2011; Gandhi et al., 2020). From here, any distortion that generates a difference between the $\mathcal{M} = \Delta = 1$ benchmark and realized input choices appears as a wedge in our model. And because we're interested in trends that are much longer than business cycles, the primary candidates are structural economic phenomena that might be evolving over decadal horizons.

On the product market side, equation 4 shows distortions $\mathcal{M}$ affect both intermediate- and labor-input choices symmetrically. Distortions in the product market add a wedge between cost shares and output elasticities for both intermediates and labor in the same multiplicative fashion. In terms of secular trends, this $\mathcal{M}$ wedge might be the result of product market power (De Loecker and Warzynski, 2012; Traina, 2018; De Loecker et al., 2020) or mismeasurement (Byrne et al., 2016). However, as we ultimately use equation 5 in practice, these trends are not confounders for our main analysis. By comparing the relative wedges of intermediates and labor, we're free to leave general potential product market distortions.

On the labor market side, equation 4 shows distortions $\Delta$ affect labor-input choices only. Although labor market power is just one possible microfoundation for this wedge (Hashemi et al., 2022), we view the candidate confounders as relatively benign in our institutional setting. For example, labor adjustment costs or dynamic contracting would be picked up in $\Delta$ in much the same way as labor market power. However, for this alternative to meaningfully confound our analysis, we'd need to believe these factors are economically significant at annual frequencies and trending meaningfully. Although we're skeptical of this belief for hourly inputs of production workers (e.g., line workers, fabricators, processors, assemblers, inspectors, and so on) we also recognize differences in beliefs on the plausibility of assumptions. For benchmarking, the literature survey in Table IV of Bloom (2009) reports the highest firm-level adjustment-cost distortion for production workers reported at 0.08. In our empirical analysis, we confirm production labor hours are not lumpy—they are not characterized by spells of inactivity as predicted by large labor adjustment costs.

In terms of other limitations, we might also consider how more sophisticated labor market models would show up in our labor market power measures. Models with dynamics may find $\Delta < 1$ at times, such as when firms pay workers less than their marginal revenue product now by promising to pay more in the future. In our setting, these dynamics should smooth out in

aggregation unless highly complex, so that our measure would capture the average $\Delta$ across production workers over a plant-year. Models with rent sharing or agency frictions may also find $\Delta < 1$, even persistently (Dobbelaere and Mairesse, 2013). We interpret our estimates of $\Delta < 1$ in this fashion, because it's most consistent with manufacturing case studies.

Finally, while operationalizing equation 5 means removing product market distortions from our measure, it also means relying more heavily on the assumption that intermediate-input markets are undistorted. We also view this assumption as relatively benign in our institutional setting, where intermediate inputs are energy and materials. However, a violation of this assumption would mean we'd mismeasure our labor market power wedge $\Delta$, possibly in a time-varying way. For example, if firms have market power over intermediates, we'd incorrectly undermeasure labor market power. More concerning, if firms receive substantial quantity discounts, we'd incorrectly overmeasure labor market power. In our empirical analysis, we confirm our results relying on equation $[L]$ of 4 and assuming $\mathcal{M} = 1$ instead; indeed, as suggested by Figure 1, product market distortions turn out to be minimal in our institutional setting.

# 2 Overview of the US Manufacturing Sector

In this section, we build on Figure (1) by documenting the trends in costs for the US manufacturing sector. In particular, we seek to understand what is driving the extraordinary rise in the cost ratio $\frac{CM}{WL}$. It's largely driven by a decline in production-worker employment $L$. The lack of trends in other variables, such as intermediate-input costs $CM$, or the pay or employment of nonproduction labor $N$, offers preliminary but strongly suggestive evidence for structural transformation in the production worker labor market. Indeed, adding to our Figure (1), we find a sudden break in $WL$ trends driven by a collapse in $L$ around the 2000 inflection point.

Figure (2) plots the time trends of the intermediate-input cost $CM$ and pro-

Figure 2: Intermediate Input Costs and Production Labor Costs

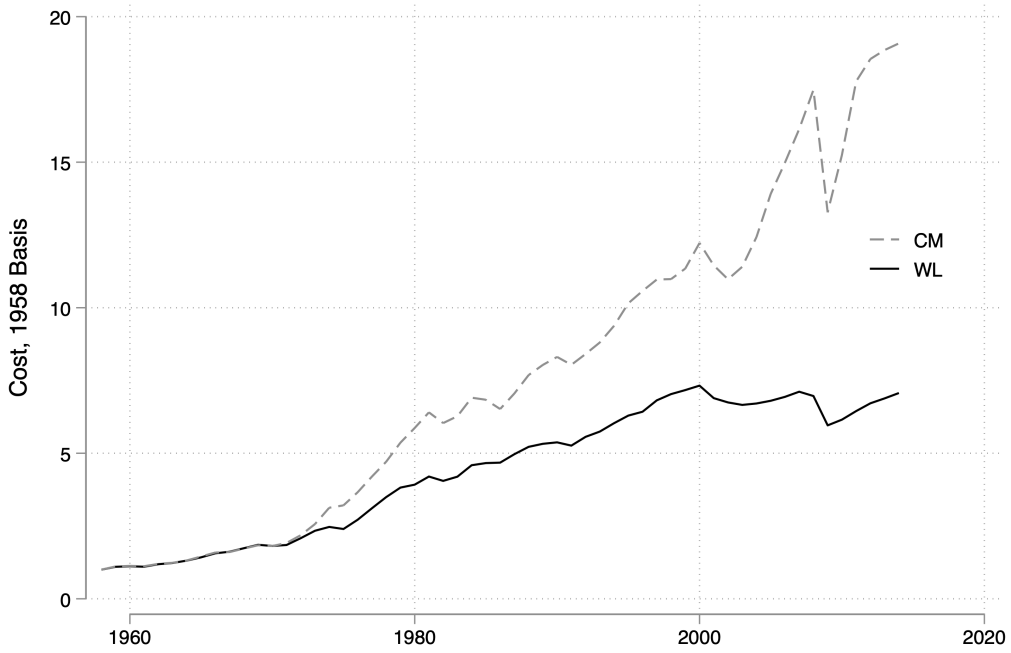


Notes: Authors' calculations using the NBER-CES Manufacturing Industry Database, March 2021 file. Intermediate input cost $CM$ is the total expenditure on energy and materials; labor-input cost $WL$ is the total expenditure on production workers. Labor costs exclude nonproduction workers such as technology professionals and managers.

duction labor cost $WL$ from 1958 to 2014. The y-axis is the cost using 1958 as a basis. From 1958 to the mid-1970s, the two lines trend upward together. However, in the mid-1970s, the $WL$ line slows down while the $CM$ line continues at pace. In 1980, $CM$ is about 6 while $WL$ is about 4. In 2000, $CM$ is about 12, while $WL$ is about 8. From 2000 onward, $CM$ continues to climb, but $WL$ remains steady.

The key conclusion from this figure is that the trends in Figure (1) are about changes in production worker markets, not intermediate input markets. This conclusion has two implications. First, it means intermediate input markets look relatively undistorted, so that we can infer levels of markdowns, not just

Figure 3: Production Labor Costs: Pay and Employment

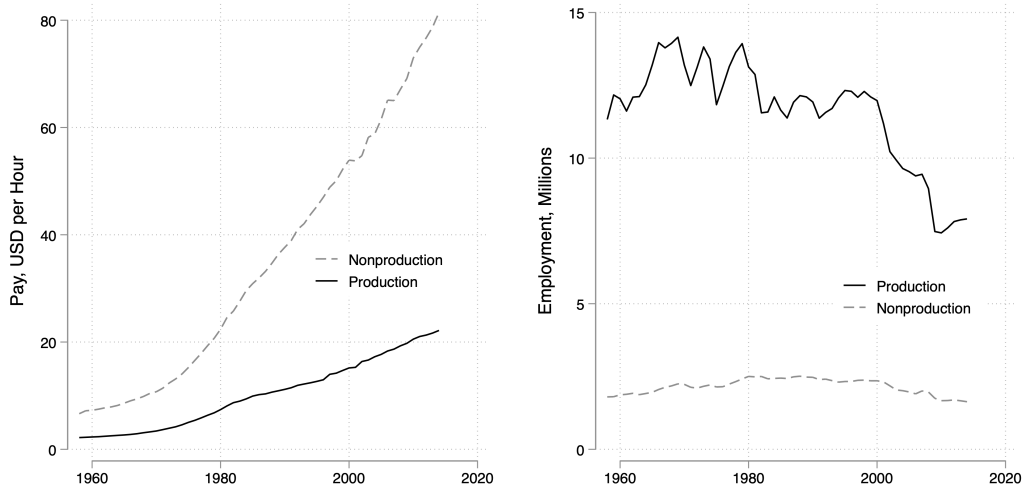


Notes: Authors' calculations using the NBER-CES Manufacturing Industry Database, March 2021 file. Production labor pay is the total expenditure on production workers divided by production worker hours. Nonproduction labor pay is the total labor expenditure bill minus total expenditure on production workers, all divided by nonproduction labor employment. Nonproduction labor employment is total employment minus production worker employment. Production labor employment in the right panel is production worker hours converted into 2000 hour equivalents.

trends. Second, it means it's unlikely we're picking up sudden changes in measurement that would affect intermediates, such as a sudden rise in offshoring.

Figure (3) contains two plots depicting trends for US manufacturing workers split by production (solid black lines) and nonproduction (dashed gray lines). The plot on the left is the trend of pay per hour from 1958 to 2014. In 1958, production pay is 2 USD per hour, and nonproduction pay is 7 USD per hour. Though both production and nonproduction pay trend upward, production pay grows slower than nonproduction pay. By the end of the panel, production pay is 21 USD per hour, and nonproduction pay is 80 USD per hour.

The plot on the right is the trend of employment. Historically, there were over 12 million production workers employed in the sector. However, 2000 marks a

breaking point when production employment sharply declines to 8 million by 2014. By contrast, during the entire panel from 1958 to 2014, nonproduction employment remains steady at about 2 million.

Taken together, the main takeaway is that the important swings in the $\frac{CM}{WL}$ cost ratio come from big drops in $L$.

# 3 Estimating Elasticities

Having established fundamental shifts in $\frac{C_{jt}M_{it}}{W_{it}L_{it}}$ cost ratios, we now turn to our microdata application that recovers $\frac{f_{l\,it}}{f_{m\,it}}$ elasticity ratios. In this section, we present our econometric approach that closely follows our companion paper (Kirov et al., 2023), adapting the production estimators of Olley and Pakes (1996), Levinsohn and Petrin (2003), and Ackerberg et al. (2015) to our institutional setting. We apply contemporary methods in production function estimation by imposing structure on the scale elasticity (Flynn et al., 2019) and latent markup determinants (Kirov et al., 2023) to overcome the nonidentification critiques of Gandhi et al. (2020) and Bond et al. (2021). Finally, we lay out our empirical specification, notably how we model plant production functions to flexibly and pragmatically allow output elasticities to change over time.

## 3.1 Empirical Strategy

In this section, we develop our control-function approach to estimating production functions in the context of revenue data, with detailed discussion in our companion paper Kirov et al. (2023). Our approach has two stages. In the first stage, we estimate the first-order condition for intermediates, in the spirit of Gandhi et al. (2020) but allowing for market power in product markets. The key idea is to use variation in inputs along with plant and year fixed effects to control for variation in markups, in the same way the proxy variable literature uses investment or intermediate input demand to control for productivity (Olley and Pakes, 1996; Levinsohn and Petrin, 2003; Ackerberg et al., 2015).

In the second stage, we combine our first-stage estimates with a Markovian revenue productivity process to identify the production function. By applying our production function estimates to firm first-order conditions, we also place ourselves in the tradition of Hall (1988) and De Loecker and Warzynski (2012) for product markets, and Dobbelaere and Mairesse (2013) and Morlacco (2020) for labor markets. Overall, our approach imposes minimal structure on supply and demand, and can recover labor and product wedges even in settings where the very nature of competition is changing.

Keeping the goal in mind, to apply our equation 5 for labor market power, we need estimates of plant production functions and their associated elasticities $f_{l_{it}}$ and $f_{m_{it}}$. The task at hand is made more complicated by the fact that we only observe revenue, not prices and quantities separately. We thus modify our earlier theoretical derivation to make it production ready for empirical implementation.

One might ask: Why do we even need advanced econometric machinery? To start, we cannot simply regress revenue on inputs to get elasticity estimates for two critical reasons. The first reason is the omitted price bias emphasized in Klette and Griliches (1996): higher markups induce plants to decrease input use, which increases prices and thus revenues. The second reason is transmission bias: higher physical productivity induces plants to increase input use, which increases output and thus revenues. At best, we'd consistently estimate revenue elasticities, but our formulas for $\mathcal{M}_{it}$ and $\Delta_{it}$ require physical output elasticities; that is, they must describe marginal increases in physical production $Q$ rather than marginal increases in revenue $R = P(Q)Q$. Although common in practice, not only would using revenue elasticities as though they were physical output elasticities generate inconsistent estimates of market power wedges (Klette and Griliches, 1996), but it'd actually recover $\mathcal{M}_{it} = 1$ under correct specification (Bond et al., 2021).[4]

[4]Our other companion paper Hashemi et al. (2022) shows revenue elasticities are actually sufficient to recover markdowns, even if they cannot recover markups. Despite this result, we do not estimate revenue elasticities in this application, as we are unaware of estimators designed to identify revenue functions. Instead, in what follows, we develop a method to

We also cannot use approaches that set output elasticities equal to cost shares (either of revenues or costs). This method would imply labor wedges equal to 1 as a tautology, since:

$$\Delta_{it} = \frac{f_{lit}}{f_{mit}} \frac{C_{jt} M_{it}}{W_{it} L_{it}} = \frac{\frac{W_{it} L_{it}}{R_{it}}}{\frac{C_{jt} M_{it}}{R}} \frac{C_{jt} M_{it}}{W_{it} L_{it}} = 1 \tag{6}$$

This result comes from the key assumption underlying the cost share approach, that markets are competitive.

Our empirical strategy is to incorporate the fact that we only observe revenue directly into our production model, and adapt existing methods to settings of imperfect competition.

We start with the same setup as before. We observe data for a panel of plants over periods $t = 1, 2, ..., T$. We omit panel subscripts and let the data take a short panel form: the number of plants grows large for a fixed $T$. For each plant, we observe revenue $R_{it} = P(Q_{it})Q_{it}$, competitively supplied intermediates $M_{it}$ with cost $C_{jt}$, production labor $L_{it}$ with cost $W(L_{it})$, nonproduction labor $N_{it}$, and capital $K_{it}$. Intermediates are inputs which plants transform directly into output, such as steel or partially finished goods. Production labor directly works to create output, while nonproduction labor supports production, such as through information technology or management. Capital is the stock of structures and equipment. $M_{it}$ and $L_{it}$ are flexible in the sense that they are both variable and static: plants may adjust them in each period after observing the realization of state variables such as productivity, and their choice has no dynamic implications. We assume $N_{it}$ and $K_{it}$ are fixed at time $t$, and $K_{it}$ follows a dynamic capital-accumulation process. Inputs generate output according to a constant returns to scale production function with Hicks-neutral productivity as before: $Q_{it} = A_{it} F_j(M_{it}, L_{it}; N_{it}, K_{it})$.

estimate quantity elasticities from revenue data. One advantage here is that if inputs and fixed effects were insufficient to control for markups, we'd still correctly recover markdowns. Overall, though, we view revenue function estimation as a promising avenue for future research.

As in Olley and Pakes (1996), the log productivity term $a_{it}$ is additively separable into a part known to the plant when making input decisions $\omega_{it}$ and an i.i.d. error term $\varepsilon_{it}$. In logs, plant production is thus $q_{it} = f_j(m_{it}, l_{it}, n_{it}, k_{it}) + \omega_{it} + \varepsilon_{it}$. Each plant uses expected output in its cost-minimization problem because it knows it must account for an as-yet-unknown portion of productivity $\varepsilon_{it}$.

The timing is as follows. First, plants inherit their nonproduction labor and capital from the previous period. They then use their expectation about their productivity $a_{it}$ conditional on the known part $\omega_{it}$, possibly along with other information, to plan markups $\mathcal{M}_{it}$ and markdowns $\Delta_{it}$. Then plants choose the corresponding flexible inputs $M_{it}$ and $L_{it}$ to implement this plan given the residual product demand and labor supply curves, technology, capital stock, and expected productivity. Finally, productivity, production, and market power terms are realized. Since expectations about productivity partially determine inputs, market power wedges depend on both productivity and conduct.

Given this setup, the plant's cost-minimization problem with time $t$ information is:

$$\begin{aligned} \min_{M_{it}, L_{it}} \quad & C_{jt} M_{it} + W(L_{it}) L_{it} \\ s.t. \quad & \mathbb{E}_{it}[A_{it}] F_j(M_{it}, L_{it}, N_{it}, K_{it}) = \bar{Q}_{it} \end{aligned} \tag{7}$$

which follows our earlier theoretical setup in equation 3, save for the new assumptions on information and timing.

We can manipulate the first-order condition for $[M]$ in the same way as equation 4. However, carrying the $\mathbb{E}_{it}[A_{it}]$ term through, we have a new unplanned productivity term $\frac{\mathbb{E}_{it}[\mathcal{E}_{it}]}{\mathcal{E}_{it}}$. In logs:

$$\mu_{it} = \log f_{mit} + r_{it} - cm_{it} + b_{it} - \varepsilon_{it} \tag{8}$$

where we define $b_{it} = \log \mathbb{E}_{it}[\mathcal{E}_{it}]$.

Let markups be a function of inputs and plant and time fixed effects $\iota_i$ and

$\tau_{jt}$:

$$\mu_{it} = \mu_j(m_{it}, l_{it}, n_{it}, k_{it}, \iota_i, \tau_{jt}) \tag{9}$$

One appealing feature is that in general models of competition, higher planned markups induce lower chosen intermediates; this suggests a straightforward way to microfound the markup control function in an input demand equation, as in Olley and Pakes (1996). If we rewrite the first-order condition equation 8 as $cm_{it} - r_{it} = \log f_{mit} - \mu_{it} + b_{it} - \varepsilon_{it}$, then the left-hand side is the intermediate input log cost share of revenue, and the $\log f_{mit} - \mu_{it}$ term on the right-hand side is the log revenue elasticity with respect to input $m_{it}$, a mix of supply and demand parameters. As productivity is Hicks-neutral, the elasticity term $f_{mit}$ is a function of inputs only: $f_{mit} = f_{mj}(m_{it}, l_{it}, n_{it}, k_{it})$. Combining the revenue elasticity terms into a single function $s_j(m_{it}, l_{it}, n_{it}, k_{it}, \iota_i, \tau_{jt}) = \log f_{mj}(m_{it}, l_{it}, n_{it}, k_{it}) - \mu_j(m_{it}, l_{it}, n_{it}, k_{it}, \iota_i, \tau_{jt})$, our first stage estimating equation becomes:

$$cm_{it} - r_{it} = s_j(m_{it}, l_{it}, n_{it}, k_{it}, \iota_i, \tau_{jt}) + b_{it} - \varepsilon_{it} \tag{10}$$

To operationalize this equation, we nonparametrically regress the intermediates share of revenues on inputs and fixed effects to get an estimate of the revenue elasticity $\hat{s}_{it} = \log \widehat{f_{mit}} - \mu_{it}$. We call this first estimating equation 10 the share regression: it defines elasticities and markups in terms of the observable share of intermediates expenditures to revenues. The share regression uses firm optimization to address the revenue problem: the left-hand side is in revenue terms, the right-hand side is in quantity terms, and the markup $\mu_{it}$ connects the two. This share regression estimates the specified markup control function: it describes the determinants of wedges between prices and marginal costs, and is similar to the share regression in Gandhi, Navarro, and Rivers (2020), but adapted for cases of imperfect competition with unobserved prices.

Importantly, it also recovers an estimate of the error $\hat{\varepsilon}_{it}$, and therefore $\hat{b}_{it}$. Estimating $\hat{\varepsilon}_{it}$ is a primary function of the first stage of proxy variable estimators

(Olley and Pakes, 1996; Levinsohn and Petrin, 2003; Ackerberg et al., 2015). Estimating it here allows us to replace the physical productivity control function assumption of these models with a markup control function assumption. However, the share regression alone cannot separate the impact of markups from output elasticities, since it still contains the unknown $f_{mit}$.

To separately identify markups and physical elasticities, we combine the share regression with structure on the revenue productivity process. Specifically, we assume revenue productivity $\nu_{it} = p_{it} + \omega_{it}$ follows a Markov process with additively separable mean-zero shocks $\eta_{it}$. Relative to the existing literature, this assumption is equivalent to the existing assumptions of Klette and Griliches (1996) and De Loecker and Warzynski (2012), and is consistent with the persistence results of Foster, Haltiwanger, and Syverson (2008).

We have:

$$\nu_{it} = g_j(\nu_{it-1}) + \eta_{it} \tag{11}$$

Following the derivation in our companion paper Kirov et al. (2023), we can combine this assumption with our first stage estimating equation to yield our second stage estimating equation, as follows.

Write firm revenues as:

$$\begin{aligned} r_{it} &= p_{it} + q_{it} \\ &= f_j(m_{it}, l_{it}, n_{it}, k_{it}) + p_{it} + \omega_{it} + \varepsilon_{it} \\ &= f_j(m_{it}, l_{it}, n_{it}, k_{it}) + \nu_{it} + \varepsilon_{it} \\ &= f_j(m_{it}, l_{it}, n_{it}, k_{it}) + g_j(\nu_{it-1}) + \eta_{it} + \varepsilon_{it} \end{aligned}$$

where the second line comes from the definition of $f_j(m_{it}, l_{it}, n_{it}, k_{it})$, the third line uses the definition of $\nu_{it}$, and the fourth line comes from our revenue-

productivity assumption. Now rewrite the output of our first stage as:

$$\begin{aligned} cm_{it} - r_{it} &= \hat{s}_{it} + \hat{b}_{it} - \hat{\varepsilon}_{it} \\ cm_{it} - f_j(m_{it}, l_{it}, n_{it}, k_{it}) - \nu_{it} - \hat{\varepsilon}_{it} &= \hat{s}_{it} + \hat{b}_{it} - \hat{\varepsilon}_{it} \\ \nu_{it} &= cm_{it} - f_j(m_{it}, l_{it}, n_{it}, k_{it}) - \hat{s}_{it} - \hat{b}_{it} \end{aligned}$$

Insert this definition into the solution for revenues to yield:

$$\begin{aligned} r_{it} =& f_j(m_{it}, l_{it}, n_{it}, k_{it}) \\ &+ g_j(cm_{it-1} - f_j(m_{it-1}, l_{it-1}, n_{it-1}, k_{it-1}) - \hat{s}_{it-1} - \hat{b}_{it-1}]) \\ &+ \eta_{it} + \hat{\varepsilon}_{it} \end{aligned}$$

Two notes are in order. First, as shown in Flynn, Traina, and Gandhi (2019), we require a scale-elasticity assumption. Second, there must be independent variation in $\mu_{it}$ which does not enter $f_j(m_{it}, l_{it}, n_{it}, k_{it})$. In our specification, this variation must come from the fixed effects $\iota_i$ and $\tau_{jt}$. With this variation, the model identifies physical quantity elasticities and markups; without this variation, the nonparametric underidentification arguments of Gandhi, Navarro, and Rivers (2020) apply. However, we view this requirement as minimal—the existing revenue productivity literature suggests persistent dispersion across plants that is independent of other latent variables such as physical productivity (Foster, Haltiwanger, and Syverson, 2008).

Our estimator is related to the proxy production function estimation model commonly used in the literature (Olley and Pakes, 1996; Levinsohn and Petrin, 2003; Ackerberg et al., 2015; Gandhi et al., 2020). We modify the proxy structure to account for the fact that we only have revenue data, and relax a key assumption in these models. In particular, we impose a Markov timing assumption on revenue productivity rather than physical productivity. Proxy models add an assumption that intermediate demand is a monotonic function of other inputs and productivity. Such a monotonicity assumption allows these models to invert productivity as a function of observed inputs. We do

not require a monotonicity assumption since we directly control for markups. This allows us to relax the necessary implication that productivity has no other determinants besides observed inputs. This scalar unobservable assumption is a requirement in such models because productivity must be inverted as a function of inputs. It is one of the more stylized assumptions in the proxy-model literature, and relaxing it is therefore valuable.

## 3.2 Production Function Estimation

Three specification decisions remain: (1) estimation groups, (2) functional forms for $f$ and $g$, and (3) instruments. We go through each in turn.

For (1) estimation groups, we pool together plants in the same 6-digit NAICS. Estimating separate production functions for narrowly defined industries is pretty standard in the literature, and to some extent we follow this precedent to tie our hands. However, we recognize that this method comes with limitations. Most notably, it does not easily accommodate structural breaks in production technologies. This possibility is ever more important as interest piques on understanding production functions over long horizons.

For (2) functional form for $f$, we use a translog production function to allow for significant flexibility in describing technology and technological change, since output elasticities depend on input choices. The four-input translog production function is:

$$
\begin{aligned}
q_{it} = \; & \theta_j^{[m]} m_{it} + \theta_j^{[l]} l_{it} + \theta_j^{[n]} n_{it} + \theta_j^{[k]} k_{it} + \\
& \theta_j^{[mm]} m_{it}^2 + \theta_j^{[ll]} l_{it}^2 + \theta_j^{[nn]} n_{it}^2 + \theta_j^{[kk]} k_{it}^2 + \\
& \theta_j^{[ml]} m_{it} l_{it} + \theta_j^{[mn]} m_{it} n_{it} + \theta_j^{[mk]} m_{it} k_{it} + \\
& \theta_j^{[ln]} l_{it} n_{it} + \theta_j^{[lk]} l_{it} k_{it} + \theta_j^{[nk]} n_{it} k_{it}
\end{aligned} \tag{12}
$$

where the $\theta$ terms are parameters estimated separately for each group $j$.

Translog production functions let our primary objects of interest $f_{m_{it}}$ and $f_{l_{it}}$

vary at the observation level $it$ through input use. To see this, note:

$$\begin{array}{rcccccccccc} f_{mit} & = & \theta_j^{[m]} & + & 2\theta_j^{[mm]} m_{it} & + & \theta_j^{[ml]} l_{it} & + & \theta_j^{[mn]} n_{it} & + & \theta_j^{[mk]} k_{it} \\ f_{lit} & = & \theta_j^{[l]} & + & 2\theta_j^{[ll]} l_{it} & + & \theta_j^{[ml]} m_{it} & + & \theta_j^{[ln]} n_{it} & + & \theta_j^{[lk]} k_{it} \end{array} \tag{13}$$

where $\theta_j^{[m]}$ and $\theta_j^{[l]}$ are the Cobb-Douglas terms, and inputs $m_{it}$, $l_{it}$, $n_{it}$, and $k_{it}$ vary at the plant-year level.

We impose shape constraints to make sure our estimates are not driven by outliers. Flynn et al. (2019) shows that $f$ must satisfy a scale-elasticity assumption; we choose constant returns to scale, as it has good empirical support in the existing literature. We also impose the neoclassical restrictions of monotonicity and concavity. For numerical stability, we constrain our estimated elasticities to be greater than the 1st percentile of cost shares in the NBER-CES Manufacturing Industry Database, which are 0.233, 0.005, 0.004, and 0.057 respectively.

For (2) functional form for $g$, we simply specify an AR(1) process. An alternative quadratic Markov specification made little difference.

For (3) instruments, we use lagged log wages, nonproduction labor, and capital, as well as their squares and interactions. We also use the lagged intermediate input cost share of revenue. Intuitively, wages instrument for production labor, fixed inputs for themselves, and the intermediate input cost share of revenue for revenue productivity.

One might ask why we do not seem to have instruments for intermediates, or why we do not use lagged intermediates or production labor themselves. Gandhi et al. (2020) shows that flexible inputs are nonparametrically collinear with productivity, and therefore have no power as instruments. Intuitively, when we see high $m$ or $l$ with high $q$, it could mean that these inputs caused the higher $q$, or simply that $\omega$ was high that period which increased input demand. However, Flynn et al. (2019) shows that a scale-elasticity assumption restores identification for problems caused by one flexible input ($m$). For the other flexible input ($l$), we rely on labor market power itself so that $w$ has

power as an instrument. Fortunately but also somewhat paradoxically, wages have more power as instruments when labor market power is high or dispersed, exactly when we'd want to best pin down the output elasticities.

Altogether, the steps to implement our estimator are:

1. Regress the intermediate input log cost share of revenues ($cm_{it} - r_{it}$) on log inputs, their squares, interactions, and plant and year fixed effects. Use the predicted residual $\hat{\varepsilon}_{it}$ to form $\hat{b}_{it} = \log \hat{\mathbb{E}}[\exp(\hat{\varepsilon}_{it})]$, and therefore $\hat{s}_{it} = \log \widehat{f_{mit}} - \mu_{it} = cm_{it} - r_{it} - \hat{b}_{it} + \hat{\varepsilon}_{it}$.

2. Specify a translog form for $f$ and an AR(1) form for $g$, along with the constant returns to scale, monotonicity, and concavity constraints.

3. Combine the estimates $\hat{s}_{it}$, $\hat{b}_{it}$, and $\hat{\varepsilon}_{it}$ with data $r_{it}$ and $cm_{it}$ and the specified functional forms for $f$ and $g$ to form the revenue productivity shock $\hat{\eta}_{it} = r_{it} - f_j(m_{it}, l_{it}, n_{it}, k_{it}) - g_j(cm_{it-1} - f_j(m_{it-1}, l_{it-1}, n_{it-1}, k_{it-1}) - \hat{s}_{it-1} - \hat{b}_{it-1}) - \hat{\varepsilon}_{it}$. This shock will be a function of $f$ and $g$ parameters.

4. Estimate the parameters of $f$ and $g$ using the moment conditions formed by:

$$\mathbb{E}[\hat{\eta}_{it} \begin{bmatrix} w_{it-1}, & n_{it-1}, & k_{it-1}, \\ w_{it-1}^2, & n_{it-1}^2, & k_{it-1}^2, \\ w_{it-1}n_{it-1}, & w_{it-1}k_{it-1}, & n_{it-1}k_{it-1}, \\ \hat{s}_{it-1}, & 1 & \end{bmatrix}] = 0 \tag{14}$$

To ease replication, we use the identity matrix for our initial GMM weight, and initialize the nonlinear system at $\theta^{[m]} = 0.5$, $\theta^{[l]} = 0.1$, $\theta^{[n]} = 0.1$, $\theta^{[k]} = 0.3$ (roughly the 50th percentile of cost shares in the NBER-CES Manufacturing Industry Database) with all square and interaction parameters at 0. We initialize the AR(1) parameter for productivity at 0.9.

# 4 The Rise of Labor Market Power

In this section, we present and discuss our main results on the rise of the labor wedge in the US manufacturing sector, particularly after 2000. We also discuss potential mismeasurement of the $\frac{CM}{WL}$ cost ratio from outsourcing and offshoring, and offer evidence these concerns are not quantitatively large enough to offset our main results. Finally, we show the overall rise in labor market power comes from an acceleration of the marginal revenue product of labor starting in the 1990s, as opposed to a structural break in wage trends.

## 4.1 Labor Market Power in US Manufacturing

Figure (4) plots the time series of the labor wedge in the US manufacturing sector from 1972 to 2014. We obtain this series by averaging our plant-level $\Delta$ estimates by year, using total employment weights. The wedge is very close to 1 before 1990, suggesting little labor market power. It begins to rise in the 1990s to approximately 1.2. Then there is an inflection point around 2000 after which $\Delta$ increases rapidly to just above 2. This increase is large: at the margin, manufacturing production workers produced output valued at approximately their wages in 1972, but worth twice as much by 2014. While the inflection point of this increase coincides with China's accession to the World Trade Organization, $\Delta$ continues to climb through the end of the sample in 2014, suggesting a fundamental transformation of the sector as highlighted by recent research (Fort et al., 2018; Charles et al., 2019).

We are cautious about interpreting the level of labor market power in Figure (4) with high precision. Our estimates assume the absence of adjustment costs, or any other characteristic, that makes a plant's production labor input choice a dynamic problem. For example, in a survey of labor adjustment-cost estimates in the literature, Bloom (2009) finds the highest distortion for production workers is 0.08. That said, our finding on the trend in labor market power is likely robust to these concerns. Plausible alternatives must explain why we see an increasingly distorted labor market relative to the intermedi-

Figure 4: Rising Wage Markdowns in US Manufacturing

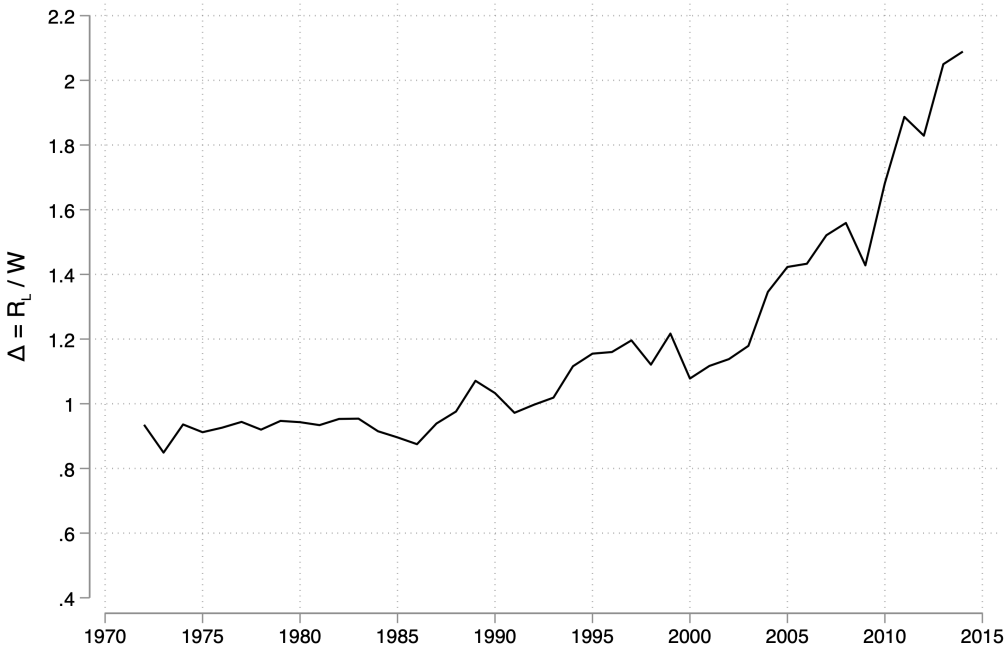


Notes: Authors' calculations using disclosed results reviewed by the US Census Bureau. Labor market power $\Delta$ is the employment-weighted average of estimated plant-level labor wedges in our US manufacturing sample. The data contain annual plant-level measures of outputs and inputs from 1972 to 2014 from our Census of Manufactures and Annual Survey of Manufactures sample.

ates market. Moreover, mismeasurement or misspecification candidates must also have inflection points at the same time and in the same direction as our findings.

One might be familiar with the work in Edmond et al. (2023) that emphasizes cost weights for markup aggregation. This prescription is only true in the particular case of all plants having the same production function, which is not true in our setting. With heterogeneous technology, there is no clear theoretical choice for wedge aggregation weights. We reason total employment is a well-measured compromise among different options.

In Figure (5), we show our estimates also imply a product wedge that is approximately 1 throughout the entire 1972 to 2014 sample. Our companion paper Hashemi et al. (2022) shows current production approach markup estimators can also pick up markdowns or other input market frictions; indeed, absent corrections for the revenue-data problem, markup estimators will measure input wedges only. Our current findings of rising markdowns and flat markups in US manufacturing suggests increases in wedges found in other settings might be driven by labor wedges rather than product wedges. For example, after adjusting the markup estimates in De Loecker et al. (2020) to account for omitted skilled labor inputs, Traina (2018) estimates a markup that rises from 1.1 to just under 1.2. However, that interpretation neglects the possibility of input market power, which would also look like a markup in the model. This problem is particularly salient given our results here and that manufacturing is overrepresented in public firm data.

The low level and stable time series of the product wedge also suggests our estimation strategy is picking up real increases in labor market power, rather than reflecting certain forms of misspecification. In particular, we could equivalently plot (4) assuming $\mathcal{M} = 1$ to find the same results of Figure (4). That is, it is sufficient to believe either intermediate input markets are undistorted, or product markets are undistorted, to conclude the labor wedge has gone up by a lot.

Figure 5: Flat Price Markups in US Manufacturing

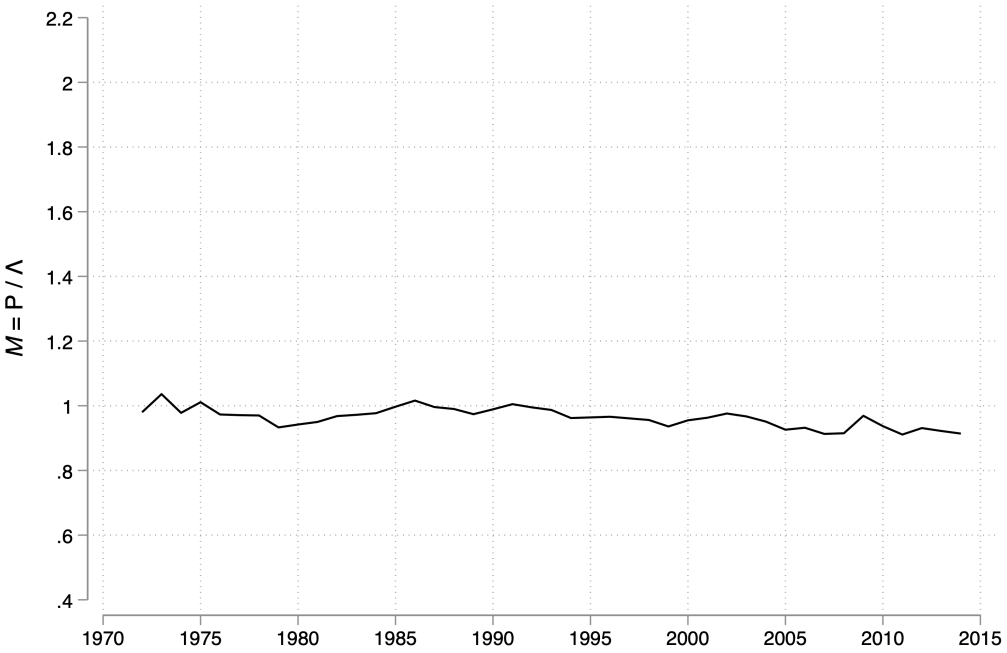


Notes: Authors' calculations using disclosed results reviewed by the US Census Bureau. Product market power $\mathcal{M}$ is the employment-weighted average of estimated plant-level product wedges in our US manufacturing sample. The data contain annual plant-level measures of outputs and inputs from 1972 to 2014 from our Census of Manufactures and Annual Survey of Manufactures sample.

Figure 6: Labor Market Power Skewness

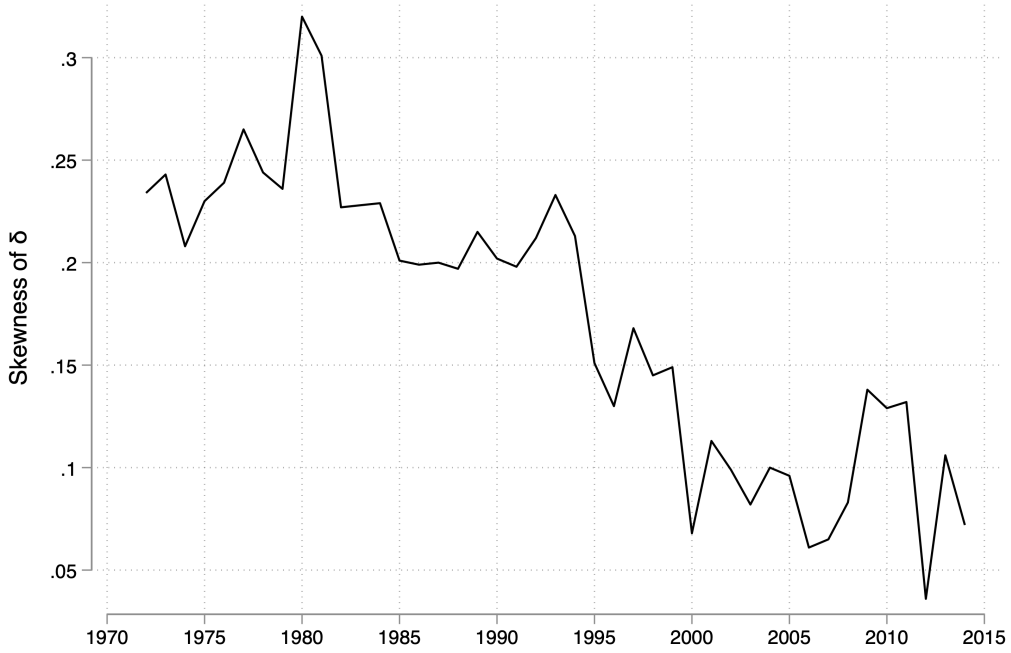


Notes: Authors' calculations using disclosed results reviewed by the US Census Bureau. Labor market power $\Delta$ is the employment-weighted skewness of estimated plant-level labor wedges in our US manufacturing sample. The data contain annual plant-level measures of outputs and inputs from 1972 to 2014 from our Census of Manufactures and Annual Survey of Manufactures sample.

Our results are also not driven by outlier plants with exceptionally high labor market power estimates. Outlier-driven results would be an especially concerning problem, as we might expect these plants are also more prone to technology misspecification or outright mismeasurement. In fact, Figure (6) shows the skewness of $\Delta$ has fallen. It plots the skewness of the log labor wedge from 1972 to 2014. Throughout the 1970s, the skewness is approximately 0.2. Then there is a steep decline from 1980 to 2000, halving to about 0.1. The trend then flattens outs.

In Figure (7), we show the time series of the cost and elasticity ratios. This elasticity ratio is simply the ratio that rationalizes the $\Delta$ series with the earlier

documented cost ratio series in Figure (1). On the left y-axis is the $\frac{CM}{WL}$ cost ratio. This trend is increasing during the panel. Beginning in 1972 at about 4, it increases modestly until 2002 to 6. After 2002, however, there is a steep and continual increase to 10 by the end of the panel in 2014. On the right y-axis is the $\frac{f_L}{f_m}$ elasticity ratio. The elasticity ratio hovers at about 0.18 with modest fluctuations from year to year. So the elasticity ratio is relatively flat, while the cost ratio rises significantly in the data. Essentially, all of our work diving into the microdata and estimating production functions only confirms the rise in the cost ratio is the story. This result is also suggestively promising for future work on US manufacturing markdowns: cost ratios are relatively accessible to researchers, and can serve as a strong benchmark for labor market power inference without the added complications of production function estimation on restricted use microdata.

We might also flip our inference and ask: What trend in the marginal product of labor would rationalize away our estimated labor market power series? Figure (8) offers an answer. We plot our estimated $f_l$ series along with the counterfactual $f_l$ series that would rationalize an undistorted production worker labor market. Here we focus on the output elasticity for labor, as the main margin that moves the cost ratio is the decline in production-worker employment. The solid line is the estimated elasticity and the dotted line is the counterfactual elasticity. The estimated elasticity is roughly constant, with modest fluctuations, including a dip down in the 1970s. The counterfactual line on the other hand shows a marked decreasing labor elasticity. It tracks the estimated elasticity until 1990, at which point it continues on to decline to 0.05 by 2014.

The estimated marginal product of labor has remained consistent at 0.10, which is also its average cost share of revenue across the sample. By contrast, the counterfactual marginal product of labor would have had to decline from 0.14 to 0.06 linearly from 1972 to 2014. While this counterfactual trend strikes us as unlikely, it does serve as an alternative view of what one would have to believe for the labor wedge to remain stable at 1.

Figure 7: The Rise of Labor Market Power: Costs and Elasticities

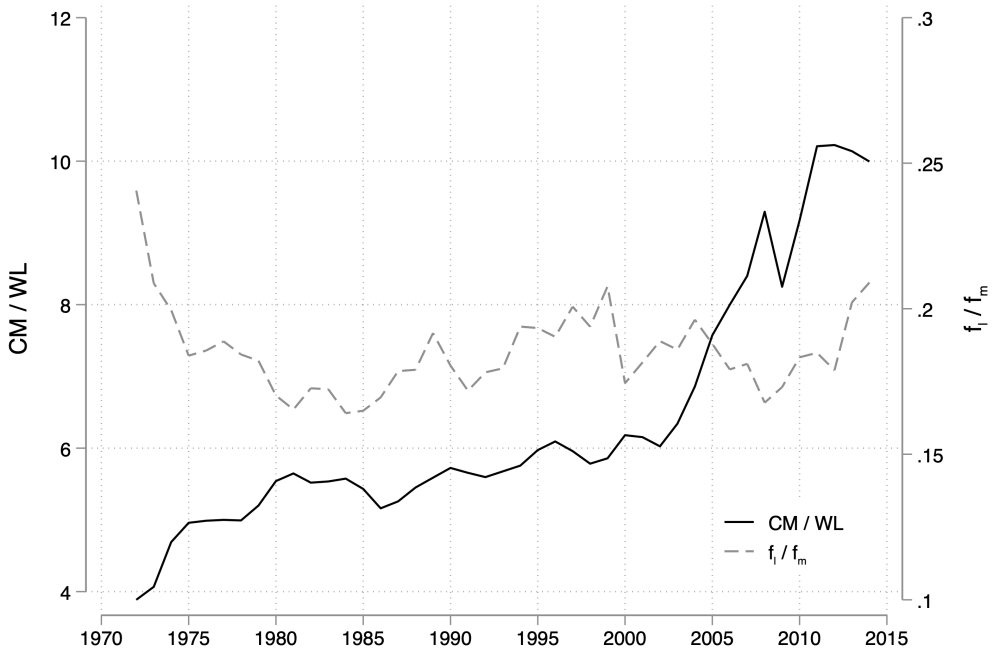


Notes: Authors' calculations using disclosed results reviewed by the US Census Bureau. The cost ratio $\frac{CM}{WL}$ is from the sector-wide totals of intermediate input and production worker expenditures. The elasticity ratio $\frac{f_l}{f_m}$ is the implied ratio that rationalizes our aggregate labor market power estimates, which are from the employment-weighted average of our estimated plant-level labor wedges. The data contain annual plant-level measures of outputs and inputs from 1972 to 2014 from our Census of Manufactures and Annual Survey of Manufactures sample.

Figure 8: The Production Labor Elasticity: Estimated and Counterfactual

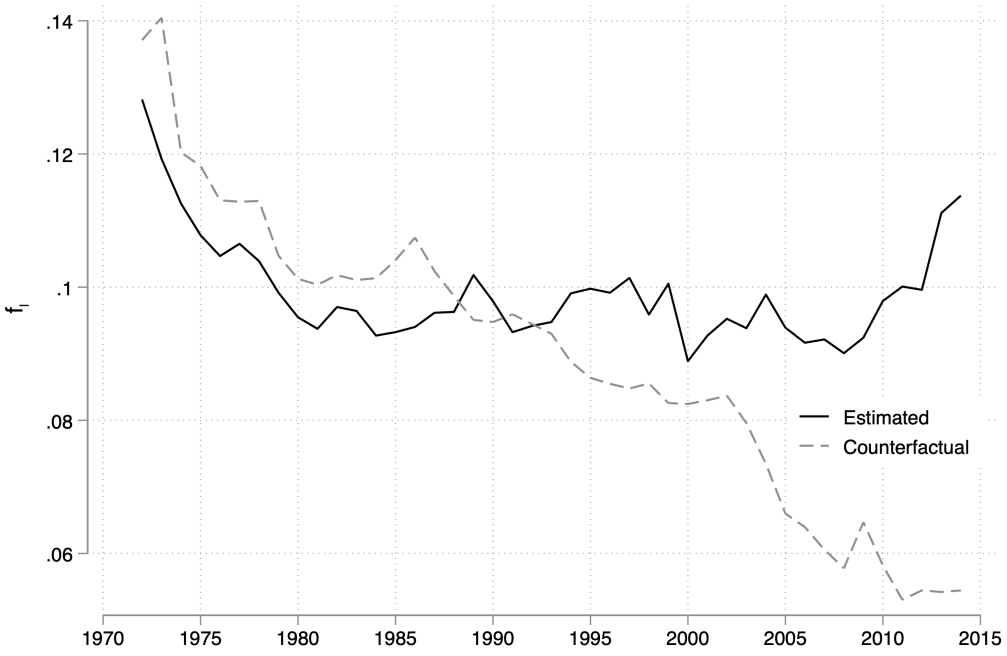


Notes: Authors' calculations using disclosed results reviewed by the US Census Bureau. The estimated output elasticity for production labor $f_l$ is the implied elasticity that rationalizes our aggregate labor market power estimates, which are from the employment-weighted average of our estimated plant-level labor wedges. The counterfactual output elasticity for production labor $f_l$ is the implied elasticity that rationalizes a hypothetical labor wedge of $\Delta = 1$ time series. The data contain annual plant-level measures of outputs and inputs from 1972 to 2014 from our Census of Manufactures and Annual Survey of Manufactures sample

## 4.2 Robustness Checks: Outsourcing and Offshoring

Having established the shifts in the $\frac{CM}{WL}$ cost ratio, and the lack thereof in the $\frac{f_l}{f_m}$ elasticity ratio, we might ask ourselves: Could our cost ratios simply be meaningfully mismeasured? In particular, we know both outsourcing and offshoring are common for US manufactures, especially over the period where we are interpreting a rise in labor market power. These relocations of economic activity might also be especially salient given the large movements in globalization and technological change. Outsourcing would mean plant-level labor bills would not include all labor truly entering the production process, while offshoring would also mean similarly omitted labor would show up in intermediate input expenditures. Both effects would mean we are overmeasuring labor market power $\Delta = \frac{f_l}{f_m}\frac{CM}{WL}$ if we were interested in the broader definition that spans outside the plant survey boundaries. (One might instead just restrict interpretation to labor market power given the measured inputs, though we view the issues as more deserving).

For outsourcing, our empirical strategy is based on the reasoning that while outsourced workers might not appear on plant labor bills, they'll appear in household surveys. So if we suppose a plant outsources a mechanic, that mechanic would leave payroll, but they'd still report they work in the manufacturing sector if asked in a census. Our strategy is to inspect how many such workers there might be, and then see what happens if we add them back into our estimates as though they were entirely omitted before.

Our main data sources to study the potential effects of outsourcing are the 1976 to 2014 Annual Social and Economic Supplements (ASEC) of the Current Population Survey, downloaded directly from the Integrated Public Use Microdata Series (IPUMS) website. We restrict the samples to prime aged workers employed in the US manufacturing sector who are not living in group quarters. We weight all data using the IPUMS ASEC individual-level survey weights.

Table (1) provides an overview of the manufacturing occupation distribution,

Table 1: The Occupation Distribution in US Manufacturing

| Occupation Group | Type | Share |
|---|---|---|
| Operators | Production | 0.32 |
| Technical, Sales, and Administrative | Nonproduction | 0.18 |
| Managerial and Professional | Nonproduction | 0.15 |
| Foremen | Production | 0.13 |
| Laborers | Production | 0.08 |
| Precision Production | Production | 0.06 |
| Repair | Outsourceable | 0.05 |
| Craft | Outsourceable | 0.02 |
| Service | Outsourceable | 0.01 |

Notes: Authors' calculations using the IPUMS Current Population Survey Annual Social and Economic Supplements, 1976 to 2014. Occupation Groups are based on the 1990 harmonized occupation codes (OCC1990) as follows: 0/200 = Managerial and Professional, 201/400 = Technical, Sales, and Administrative, 401/470 = Service, 501/550 = Repair, 551/620 = Craft, 621/700 = Precision Production, 701/800 = Operators, and 801/900 = Laborers. The exception are Foremen, which are direct supervisors defined as OCC1990 codes of 22, 503, 558, 628, or 803. All other occupation codes (Farming, Forestry, Fishing, Military, and Unclassified) form less than 1% of the sample and are dropped. The data contain demographic and employment information. As noted on the website, they must be used for good, never for evil.

pooled over the 1976 to 2014 samples. We classify the harmonized occupation codes into occupation groups in the first column, and further classify them into three buckets based on how they might appear in plant surveys in the second column. These buckets are measured production labor ("Production"), nonproduction labor ("Nonproduction"), and potentially unmeasured production labor ("Outsourceable"). We can think of the last bucket as service jobs in manufacturing that might leave the boundaries of the plant. In total, they make up anywhere from 1% to 8% of US manufacturing employment.

Figure (9) plots trends in the service share of manufacturing employment from 1976 to 2014. The lines are additive and vary based on our classification of jobs would most likely to leave labor bills. Regardless of classification, however, the service share trend is flat. If anything, it actually stabilizes in the 1990s from

Figure 9: Trends in the Service Share

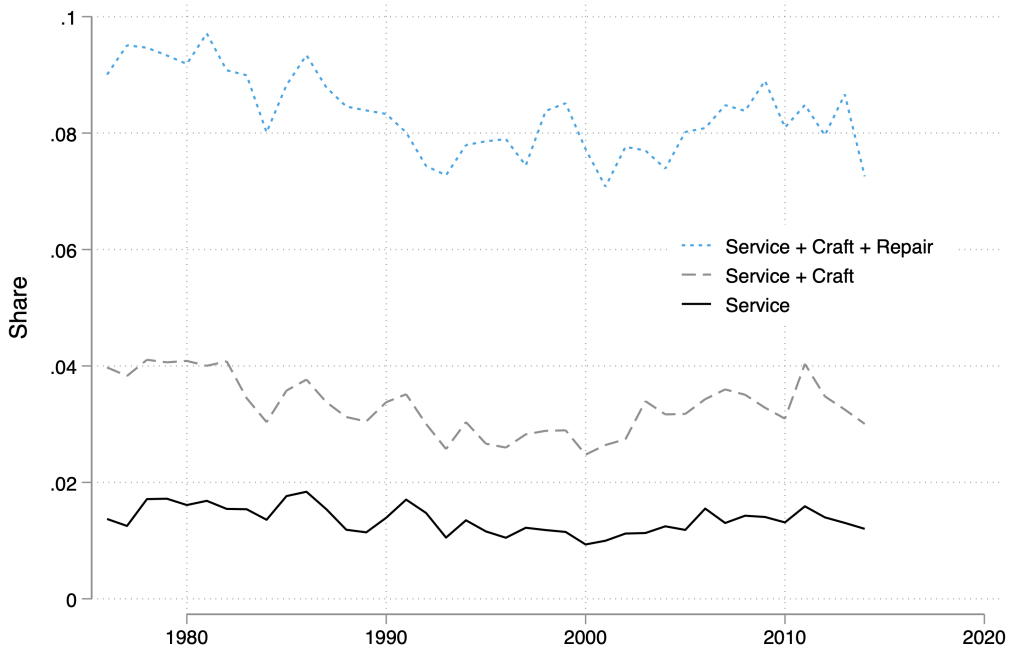


Notes: Authors' calculations using the IPUMS Current Population Survey Annual Social and Economic Supplements, 1976 to 2014. Occupation Groups are based on the 1990 harmonized occupation codes (OCC1990) as follows: 401/470 = Service, 501/550 = Repair, 551/620 = Craft, except OCC1990 codes 503 and 558 which are Foremen. All other occupation codes (Farming, Forestry, Fishing, Military, and Unclassified) form less than 1% of the sample and are dropped. The data contain demographic and employment information. As noted on the website, they must be used for good, never for evil.

a slight decline beforehand.

Figure (10) plots the rise of labor market power with outsource adjustments from 1972 to 2014 using the most conservative case, i.e., reallocating all service, craft, and repair jobs. The solid black line tracks the baseline estimates. The dashed gray line adjusts these estimates by mechanically adding back these workers. The dotted blue line adjusts these estimates further through a knock-on $f_l$ effect—intuitively, if you omit half the workers in your production function estimation, the remaining ones will look twice as productive. Even in this case, while the levels are about 20% lower than the baseline, the trends

Figure 10: The Rise of Labor Market Power, Outsource-Adjusted

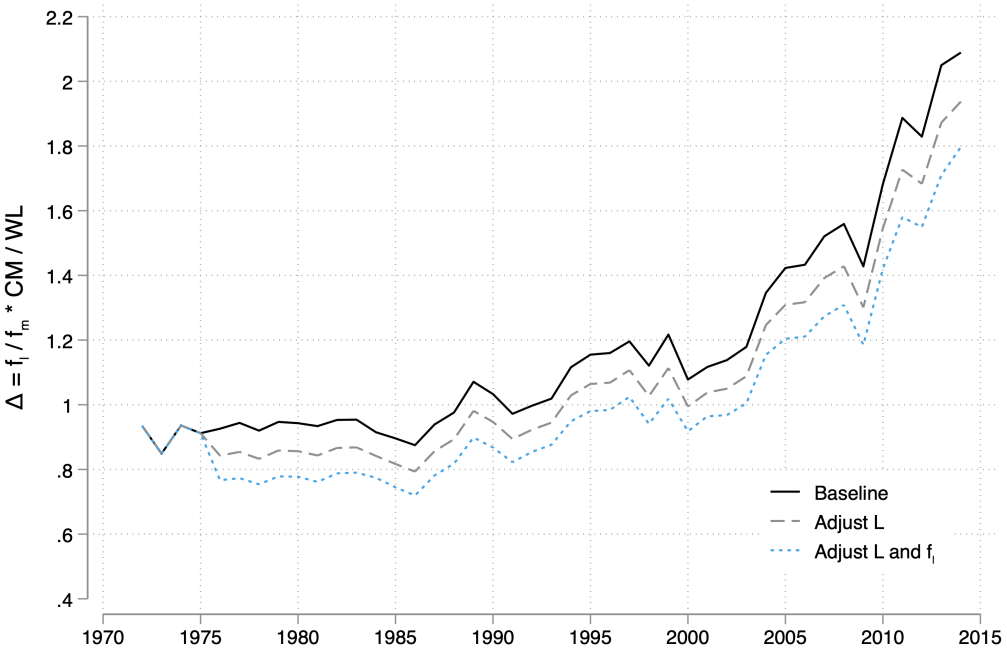


Notes: Authors' calculations using disclosed results reviewed by the US Census Bureau, and data from the IPUMS Current Population Survey Annual Social and Economic Supplements, 1976 to 2014. Labor market power $\Delta$ is the employment-weighted average of estimated plant-level labor wedges in our US manufacturing sample. Occupation Groups are based on the 1990 harmonized occupation codes (OCC1990) as follows: 401/470 = Service, 501/550 = Repair, 551/620 = Craft, except OCC1990 codes 503 and 558 which are Foremen.

remain.

Offshoring is a bit more difficult—data are scarce, much to the detriment of trade economics. Our empirical strategy is based on a similar reasoning as outsourcing: offshored workers might not appear on plant labor bills, but their productive inputs will appear as imported intermediates. So if we suppose an automotive plant offshores body production, the workers costs would leave payroll, but reappear as intermediate input costs when the body is imported back into the plant's production process. We therefore focus on import intensity, both across industries and through time. Though moving all imported intermediates from the intermediates bill to the labor bill would be a gross

overestimate, we might still believe variation in intermediate import intensity is useful when normalized, so we proceed in kind.

Our main data sources to study the potential effects of offshoring are the 1997 to 2014 BEA's Import Matrices, from the Input-Output Accounts section of the BEA website. We use the Use of Imported Commodities by Industry tables, 71 Industries.

Figure (11) plots the import intensity against the log cost ratio. Even as the import intensity increases from 0 to 0.3, the log cost ratio looks fairly constant at about 2.

Figure (12) plots the rise of labor market power with offshore adjustments from 1972 to 2014. We normalize the offshoring adjustment to the start of our BEA data in 1997, and track the relative growth from there. The two lines diverge a few years later. The baseline hits 2.1 by the end of the panel; the adjusted line hits 1.8. While offshoring might explain some rise in labor market power, it is simply too small to be the lion's share of our findings.

## 4.3 Productivity and Pay

In what sense has the markdown risen? Figure (13) shows a time series of the sector-wide average marginal revenue product of labor $R_L$ and the average wage $W$ from 1972 to 2014. From 1972 to about 1990 the productivity and pay trends are nearly equal, trending from about 4 to about 11. The average $R_L$ starts rising faster than the average wage in the 1990s. This process accelerates significantly in the 2000s, growing to about 46 USD per hour by 2014. In contrast, wages only grow to about 21 USD per hour by 2014. The labor wedge therefore increases because "productivity" rises, and not because pay falls. This suggests technological change plays a large role in the rise of the labor wedge.

Figure 11: Labor Market Power and Intermediate Imports

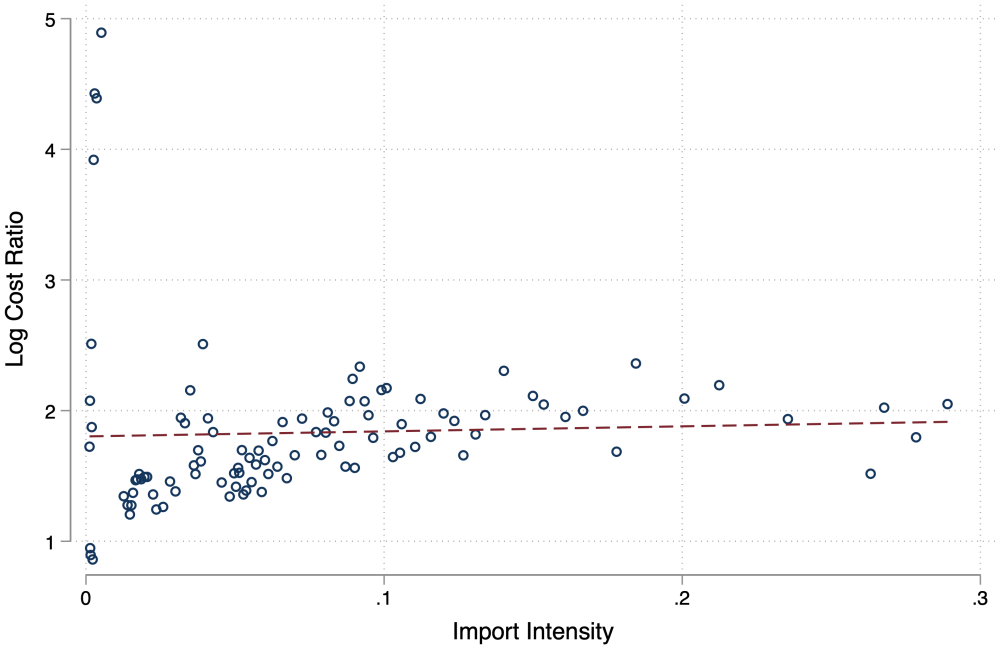


Notes: Authors' calculations using the BEA Input-Output Accounts Import Matrices, 1997 to 2014 and the NBER-CES Manufacturing Industry Database, March 2021 file. The data are at the 3-digit NAICS-year level. The Log Cost Ratio is the log of the $\frac{CM}{WL}$ cost ratio, where intermediate input cost $CM$ is the total expenditure on energy and materials and labor input cost $WL$ is the total expenditure on production workers. Import Intensity is the share of $CM$ that are intermediate imports. Intermediate imports are the total commodities imported in the same 3-digit NAICS industry, for each 3-digit NAICS industry (i.e., the diagonal terms of the Import Matrices). The BEA data show input-output production relationships among industries for 71 3-digit NAICS industries, 18 of which are in manufacturing.

Figure 12: The Rise of Labor Market Power, Offshore-Adjusted

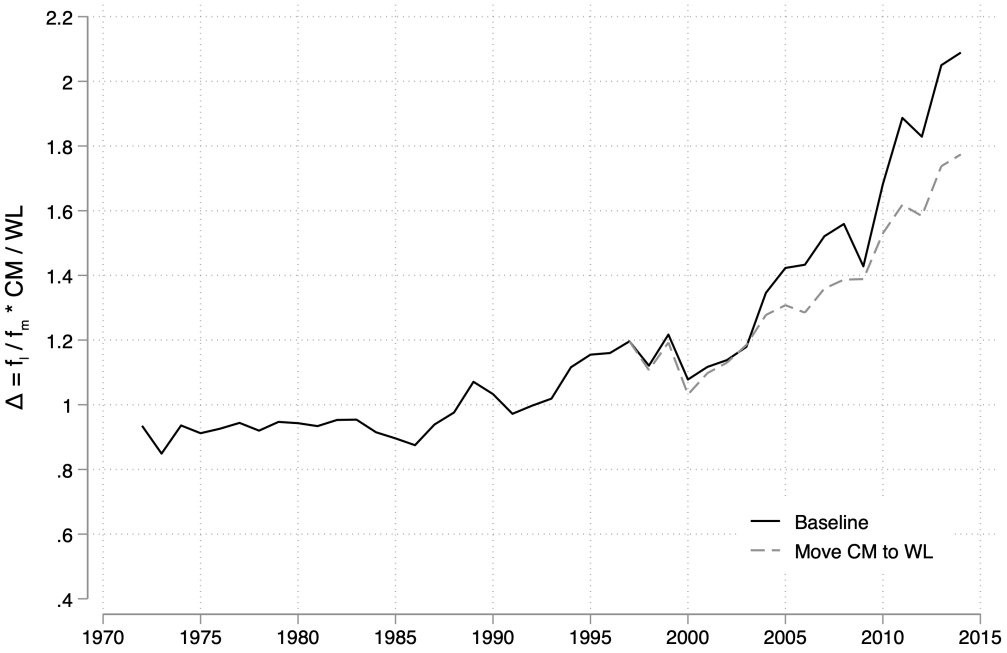


Notes: Authors' calculations using disclosed results reviewed by the US Census Bureau, and data from the BEA Input-Output Accounts Import Matrices, 1997 to 2014. Labor market power $\Delta$ is the employment-weighted average of estimated plant-level labor wedges in our US manufacturing sample. Intermediate imports are the total commodities imported in the same 3-digit NAICS industry, for each 3-digit NAICS industry (i.e., the diagonal terms of the matrices).

Figure 13: The Rise of Labor Market Power: Productivity and Pay

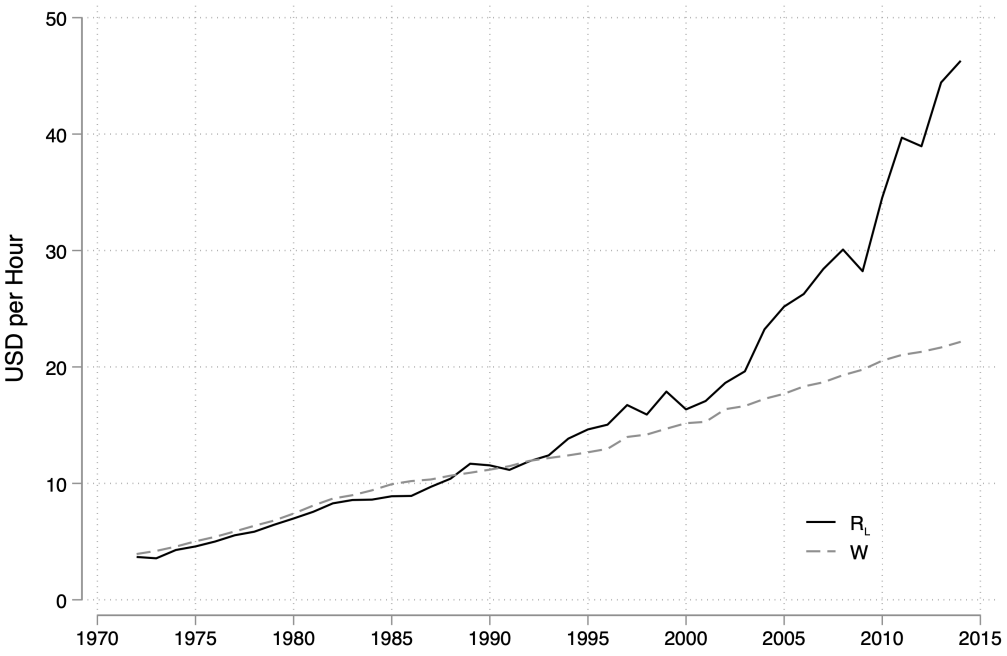


Notes: Authors' calculations using disclosed results reviewed by the US Census Bureau. The marginal revenue production of labor $R_L$ is the implied value that rationalizes our labor market power estimates. Labor market power $\Delta$ is the employment-weighted average of estimated plant-level labor wedges in our US manufacturing sample. The data contain annual plant-level measures of outputs and inputs from 1972 to 2014 from our Census of Manufactures and Annual Survey of Manufactures sample.

# 5 Implications for Labor Markets

We next investigate how labor market power relates to labor market outcomes. We begin by looking at consequences, understanding how much our wedge estimates might explain the fall in the manufacturing labor share. We then turn to market predictors of labor market power, with special attention to employer concentration because of its significant role in the emerging literature (Rinz, 2022; Berger et al., 2022; Hershbein et al., 2019; Azar et al., 2022; Benmelech et al., 2022). Many models either directly or indirectly use employer concentration as measures of labor market power, often intertwined with mega- (large) or superstar (productive) firm assumptions.

## 5.1 Labor Shares, Wages, and Employment

What proportion of the decrease in the manufacturing labor share can we plausibly explain with the increase in the labor wedge? Figure (14) shows that labor shares are negatively correlated with labor wedges. On the x-axis is the log labor wedge, and on the y-axis is the labor share of value-added. A simple regression of plant-year labor shares on $\delta = \log \Delta$ yields a coefficient of approximately -0.1: a 10% increase in $\Delta$ yields a reduction in the labor share of 1 percentage points.

Applying this estimate to the time series in a back-of-the-envelope calculation suggests that the doubling in $\Delta$ between 1972 and 2014 reduced the manufacturing labor share by 10 percentage points. The labor share in manufacturing fell by about 20 percentage points, so this calculation implies that roughly half of the decline in the manufacturing labor share might be attributable to rising labor market power. While this calculation ignores important general equilibrium effects such as the reallocation of labor and markdowns across plants over time, it nevertheless suggests that changes in labor market power are important contributors to the decline in the labor share in manufacturing.

We next move to decomposing this relationship into wages, labor inputs, and labor expenditures. We regress each in turn on our log labor wedge $\delta$ estimates

Figure 14: Labor Market Power and Labor Shares

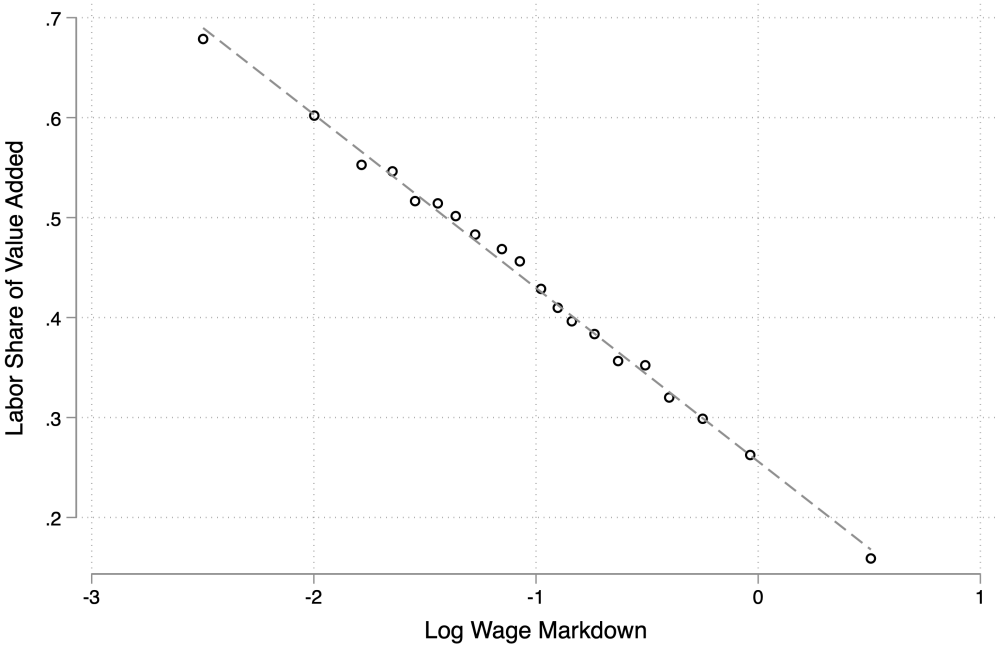


Notes: Authors' calculations using disclosed results reviewed by the US Census Bureau. Log labor market power $\delta$ is the log of our estimated plant-level labor wedges $\Delta$ in our US manufacturing sample. We weight all analyses by total employment. The data contain annual plant-level measures of outputs and inputs from 1972 to 2014 from our Census of Manufactures and Annual Survey of Manufactures sample.

Table 2: Labor Shares: Wages and Employment

| | (1) | (2) | (3) |
|---|---|---|---|
| | $w$ | $l$ | $wl$ |
| $\delta$ | $-0.027$ | $-0.173$ | $-0.193$ |
| | (0.001) | (0.002) | (0.002) |
| $R^2$ | 0.22 | 0.04 | 0.34 |
| $AIC$ | 707 | 2102 | 2189 |

Notes: Authors' calculations using disclosed results reviewed by the US Census Bureau. Log labor market power $\delta$ is the log of our estimated plant-level labor wedges $\Delta$ in our US manufacturing sample. Log labor input cost $wl$ is the log of total expenditure on production workers. We decompose this term into log hourly pay $w$ and log hours $l$. We weight all analyses by total employment. The data contain annual plant-level measures of outputs and inputs from 1972 to 2014 from our Census of Manufactures and Annual Survey of Manufactures sample.

at the plant-year level, each with 6-digit NAICS industry and year fixed effects. Table (2) collects the results.

A larger $\delta$ reduces labor inputs about six times as much as it reduces wages. Table (2) therefore shows that our labor wedge predicts quantity restrictions on the part of firms, rather than wage reductions (of course, the two are tightly linked by the firm's residual labor supply curve). This result is unsurprising in light of Figure (3), as we saw that most of the movement in labor bills comes from a drop in employment.

## 5.2 Employer Concentration and Size

Do more concentrated labor markets imply greater $\Delta$? Simple models of labor competition, such as oligopsony models that are in essence "inverted" oligopoly models, often imply greater markdowns for firms that are large in their local labor markets. We uniquely have access to measures of labor market power that do not rely on these strong assumptions on the nature of labor market competition. Hence, we can use our estimates to test these models.

In our sample of US manufacturing plants from 1972 to 2014, we find that

employer concentration and wage markdowns are only weakly related. As a measure of concentration, we compute a Herfindahl-Hirschman Index (HHI) as the sum of the squared share of production workers within each commuting zone-year. We use a commuting zone as an approximate geographic measure of a local labor market, which is fairly standard. More critically, we are limited by the fact that we do not know the relevant labor market definition along the worker dimension. Defining the relevant market for these types of antitrust questions is notoriously difficult. For example, we might rather want to include construction employers as possible outside options. We proceed with this caveat, noting that the same critique applies throughout the aforementioned studies also studying employer concentration.

We aggregate our labor wedges to the commuting zone-year level, regress $\delta = \log \Delta$ on our employer concentration measure, and display the results in Figure (15). On the left is the plot of log labor wedges on HHI in levels; on the right is the plot of log labor wedges on HHI in changes.

In the left panel, we find evidence of a positive cross-sectional relationship between labor HHIs and $\delta$. However, such a result might be driven by persistent unobservable characteristics of commuting zones, such as income or population density. To control for these confounders, we estimate the same model in changes and display the results in the right panel. Changes in a commuting zone's HHI are approximately uncorrelated with changes in its average labor wedge. Thus, labor markets which become more concentrated do not increase their labor wedges, on average.

These results suggest that the workhorse assumptions behind some of the labor market power literature might need reevaluation, particularly work that uses cross-sectional variation to infer trends in labor market power. Concentration is likely an inappropriate measure of labor market power in this case. On the bright side, these results also suggest promising avenues of future research: What is the nature of labor market competition if not static oligopsony? We explore one possibility later in the paper: technological change.

Figure 15: Employer Concentration?

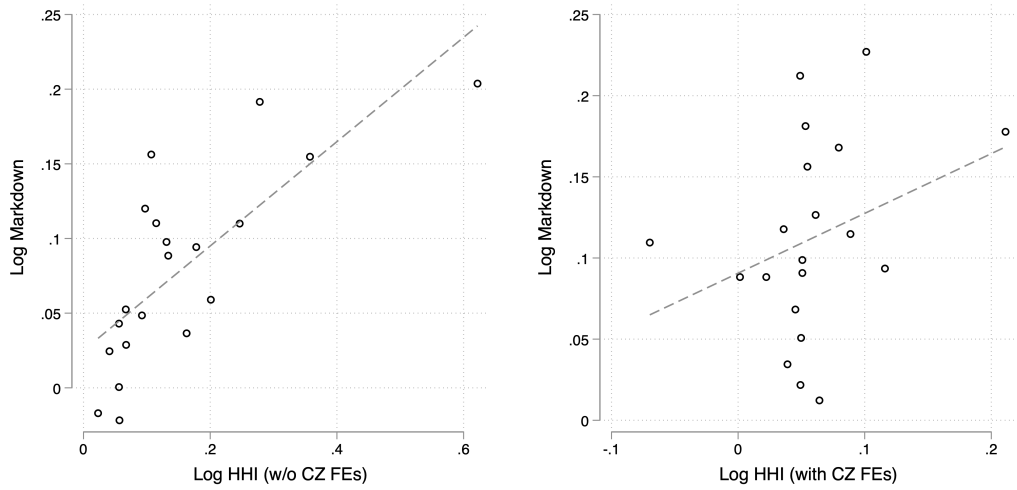


Notes: Authors' calculations using disclosed results reviewed by the US Census Bureau. Log labor market power $\delta$ is the log of our estimated plant-level labor wedges $\Delta$ in our US manufacturing sample. We weight all analyses by total employment. The data contain annual plant-level measures of outputs and inputs from 1972 to 2014 from our Census of Manufactures and Annual Survey of Manufactures sample.

Figure 16: Labor Market Power and Size?

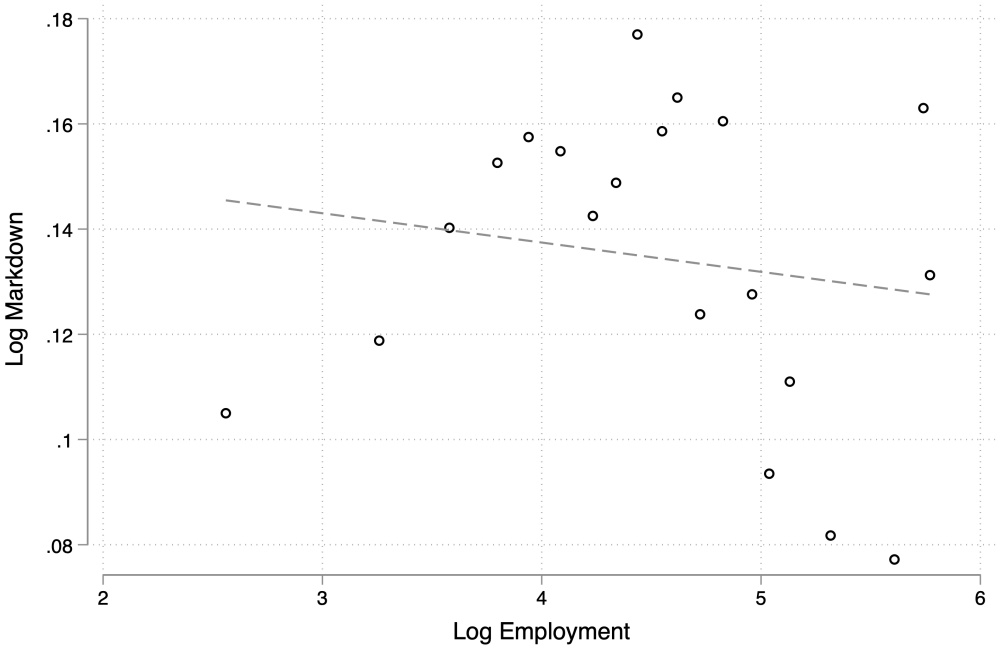


Notes: Authors' calculations using disclosed results reviewed by the US Census Bureau. Log labor market power $\delta$ is the log of our estimated plant-level labor wedges $\Delta$ in our US manufacturing sample. We weight all analyses by total employment. The data contain annual plant-level measures of outputs and inputs from 1972 to 2014 from our Census of Manufactures and Annual Survey of Manufactures sample.

As further evidence against concentration hypotheses, Figure (16) examines the relationship between the log labor wedge $\delta$ and size. We plot the binscatter of our labor market power estimates against the log of total employment (i.e., including nonproduction workers), controlling for plant and year fixed effects. Plant-years with larger labor wedges are not systematically bigger. Combined with the evidence in Table (2), this finding suggests that plants with high labor wedges reduce production employment and increase nonproduction employment: large $\delta$ are associated with more intensive use of managerial technologies.

We turn next to some other suggestive evidence on the sources of labor market

Table 3: Analysis of Variance: Industries and Commuting Zones

| | (1) $\delta$ | (2) $\delta$ | (3) $\delta$ | (4) $\delta$ |
|---|---|---|---|---|
| Year FEs | X | X | X | X |
| Industry FEs | | X | | X |
| Commuting Zone FEs | | | X | X |
| $R^2$ | 0.01 | 0.81 | 0.11 | 0.82 |
| $AIC$ | 2343 | 1334 | 2277 | 1302 |

Notes: Authors' calculations using disclosed results reviewed by the US Census Bureau. The data contain annual plant-level measures of outputs and inputs from 1972 to 2014 from our Census of Manufactures and Annual Survey of Manufactures sample.

power. We ask: Is labor market power driven more by technology or institutions? To get at this question, Table (3) regresses plant-year log labor wedges on different sets of fixed effects. The models we run are of the form:

$$\delta_{jct} = \alpha_j + \gamma_c + \tau_t + \varepsilon_{jct}$$

where $j$ is a 6-digit NAICS industry, $c$ is a commuting zone, and $t$ is a year. We weight each regression by total employment.

Column 1 of Table (3) includes only time fixed effects, and has a low R-squared and high AIC: it is not a particularly predictive model. Column 2 adds 6-digit NAICS industry fixed effects (there are 364 such codes). The R-squared of the regression increases dramatically from 0.01 to 0.81, and the AIC falls by almost half. Thus, industry adds significant explanatory power. Columns 3 and 4 look at whether a similar statement can be said for commuting zones, which are directly relevant for employer concentration. The commuting zone and year fixed effects model in Column 3 only has an R-squared of 0.11, while adding industry fixed effects increases it eightfold and reduces the AIC by nearly half. In other words, the rise in labor market power is explained much more by a plant's industry than by its commuting zone. This suggests that changes in technology are an important contributor to increased sector-wide labor market

power: technological change at the industry level explains $\delta$ better than any institutional change at the geographic level.

In sum, plants with higher labor wedges reduce hiring more than they reduce pay, while substituting production worker for managers. They also have significantly smaller labor shares of value-added: indeed, a rough calculation suggests that half of the decline in the US manufacturing labor share can be attributed to increased labor market power. If not employer concentration, what drives increases in labor market power? We turn to this question next.

# 6 The Role of Technological Change

Figure (4) shows that the labor wedge has inflection points upward in the early 1990s and the early 2000s. These events coincide with important changes in technology, and suggest a technological explanation for the labor wedge. In this section, we directly examine how a plants' technology predicts its labor wedge.

## 6.1 Labor Market Power and Technological Change

The manufacturing microdata include several measures of information and communication-related technological change in recent years. In 2000 and 2001, plants were asked to report their new computer expenditures. And in 1997, 2002, and each year after 2006, they were asked about their cost of purchased communications. We normalize each of these measures by dividing by total plant employment, then regress log labor wedges on these technological ratios while controlling for industry and year fixed effects.[5]

Figure (17) shows cross-sectional binscatters of our regression results for infor-

[5]We normalize by the number of workers (production workers plus nonproduction workers) to obtain a measure of technological intensity, rather than pick up effects from plant size. The idea is to test whether plants which spend more on technology relative to their workforce differ systematically in labor market power. Our results are robust to normalizing these measures by revenues instead of number of workers, as well as to directly controlling for plant size.

Figure 17: Information and Communication Technologies

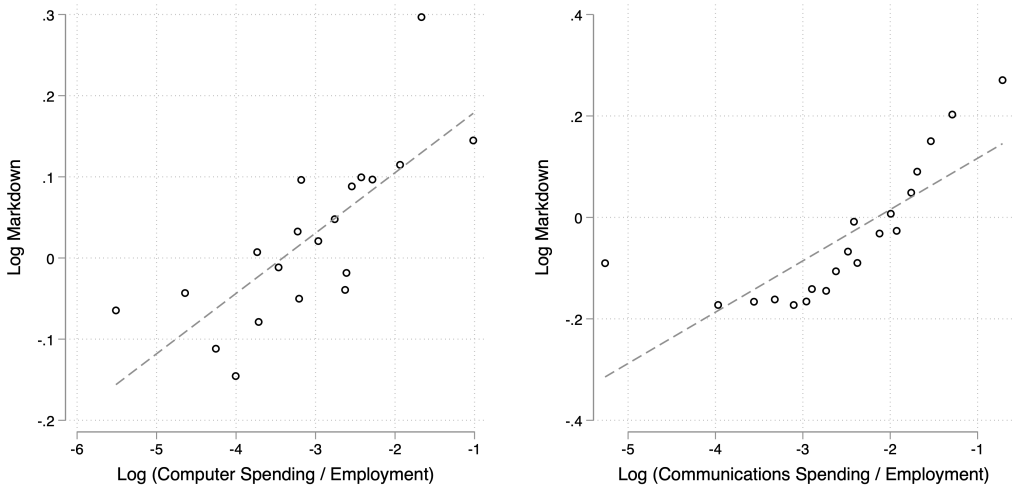


Notes: Authors' calculations using disclosed results reviewed by the US Census Bureau. Log labor market power $\delta$ is the log of our estimated plant-level labor wedges $\Delta$ in our US manufacturing sample. We weight all analyses by total employment. The data contain annual plant-level measures of outputs and inputs from 1972 to 2014 from our Census of Manufactures and Annual Survey of Manufactures sample.

mation and communication technology and the log labor wedge. Both graphs have log labor wedge on the y-axis. The graph on the left has the log of new computer expenditures per worker on the x-axis, while the graph on the right has the log of new computers expenditures per worker on the x-axis. Both technology measures are strongly predictive.

In each case, plants which spend more on information and communication technology have larger labor wedges. This pattern is consistent across both technology spending measures. Thus, plants which buy more computers and more communications have higher labor wedges.

We see the results in this section as strong evidence for a technology-centered explanation of the increase in the sector-wide labor wedge. We've confirmed that these relationships are not driven by outliers, and hold with more granular panel fixed effects (firm and plant). In unreported results, we also found that technological spending noisily predicts a plant's position in the distribution

of labor wedges in subsequent years, particularly for computer spending. In short, technological investment correlates strongly with labor wedges.

In addition to the direct measures of technological intensity, plants' labor wedges are positively correlated with measures of technology-intensive inputs. In particular, we examine whether plants which more intensively use nonproduction labor or capital have higher labor wedges. We interpret these inputs as proxies for managerial and automation technologies (See also Atalay et al. (2014)).

We regress plant-level labor wedges on nonproduction intensity and capital intensity, measured as the log of the ratio of $N$ or $K$ to total employment. As with the earlier analyses, we include 6-digit NAICS and year fixed effects. Figure (18) plots the automation and management technologies and the log labor wedge. In the graph on the left, the x-axis is the log of the number of nonproduction workers over the total number of workers. There is a positive relationship with log labor wedges. In the graph on the right, the x-axis is log capital over number of workers. Again there is a notable positive relationship.

In sum, we find that plants with greater technology intensity likely also have higher labor wedges. This section thus provides robust supporting evidence for technology playing a key role in the nature of labor market competition in US manufacturing.

## 6.2 Supporting Evidence: Unions and Offshorability

Finally, we turn to testing other determinants of labor market power, for which data are a bit more scarce. The rise in the wedge coincides with a fall in manufacturing unionization and significant changes to manufacturing technology. How important are these factors? We investigate this question using data from the Current Population Survey, as well as information on job tasks from Autor and Dorn (2013). We consider the fraction of prime aged workers in the labor force which are unionized as a potential driver of labor wedges, as well as job offshorability. We see these explanatory variables

Figure 18: Automation and Management Technologies

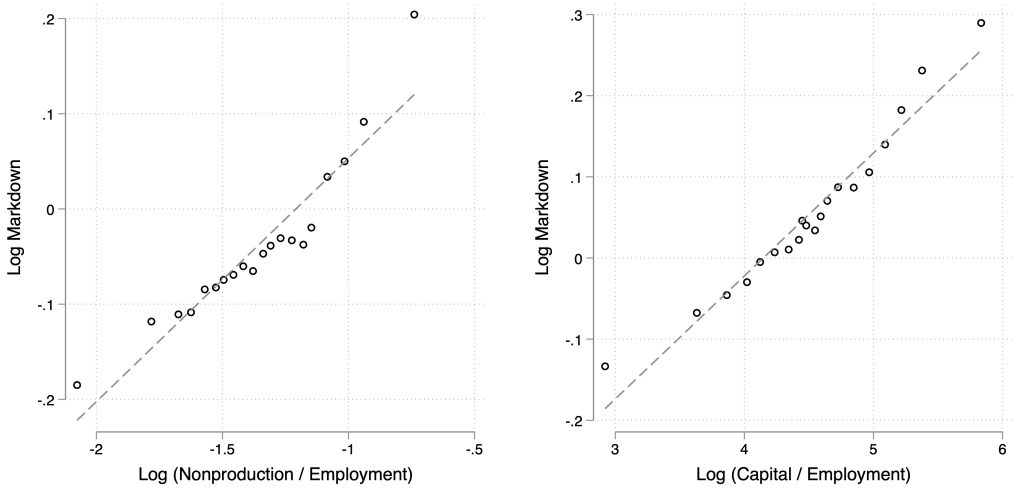


Notes: Authors' calculations using disclosed results reviewed by the US Census Bureau. Log labor market power $\delta$ is the log of our estimated plant-level labor wedges $\Delta$ in our US manufacturing sample. We weight all analyses by total employment. The data contain annual plant-level measures of outputs and inputs from 1972 to 2014 from our Census of Manufactures and Annual Survey of Manufactures sample.

Table 4: Unions and Offshorability

| | Industries | | Commuting Zones | |
|---|---|---|---|---|
| | (1) | (2) | (3) | (4) |
| | $\delta_{2010-1990}$ | $\delta_{2010-1990}$ | $\delta_{2010-1990}$ | $\delta_{2010-1990}$ |
| $\text{Union}_{2010-1990}$ | −1.65 | | −0.18 | |
| | (0.56) | | (0.72) | |
| $\text{Offshorable}_{2010-1990}$ | | −1.04 | | −0.91 |
| | | (0.89) | | (1.11) |
| *Obs* | 70 | 70 | 550 | 550 |
| $R^2$ | 0.12 | 0.02 | 0.00 | 0.00 |
| *AIC* | 102 | 113 | 1148 | 1147 |

Notes: Authors' calculations using disclosed results reviewed by the US Census Bureau, and the IPUMS Current Population Survey. Log labor market power $\delta$ is the log of our estimated plant-level labor wedges $\Delta$ in our US manufacturing sample. We weight all analyses by total employment. Union rates come from the IPUMS data directly. Offshorable rates come from merging the IPUMS 1990 harmonized occupation codes (OCC1990) with the offshorability measure in Autor and Dorn (2013), downloaded from the David Dorn Data Page.

as broadly reflecting institutional and technological changes in labor market conditions. We run analyses at both the subsector (roughly 3-digit NAICS) and commuting zone levels.

Table (4) shows the results, focusing on the rise from 1990 to 2010. Columns 1 and 2 are in long differences at the subsector level, and Columns 3 and 4 are in long differences at the commuting zone level. There are two main takeaways. First, industries with greater union membership have lower labor wedges, suggesting a mediating role. This result does not hold for commuting zones, perhaps unsurprisingly, given that geography explains little of the variation in Table (3). Second, confirming the offshoring analysis of Figure (12), the offshorability of jobs in either an industry or commuting zone displays little relation.

# 7 Concluding Remarks

In this paper, we find labor market power over US manufacturing production workers rose substantially. Firms paid production workers in the manufacturing sector their marginal revenue product in 1972, but only half this amount by 2014, showing a growing labor wedge. This wedge emerges because marginal-revenue-product growth speeds up while wage growth remains stable. The rise in the labor wedge was sharp in the early 2000s. At the plant level, high wage markdowns predict low labor shares, consistent with the hypothesis that labor market power helps account for the decline in the US manufacturing labor share.

Our results underscore technological change as an essential driver of labor market power. We find mixed evidence relating wage markdowns to employer concentration. At the industry level, increases in markdowns correlate with decreases in unionization. At the plant level, markdowns strongly correlate with technology-related expenditures on computers and communications, in both levels and changes. Markdowns also strongly correlate with indirect measures of managerial and automation technologies, as proxied by nonproduction labor and capital intensities.

Overall, our paper has two primary conclusions. First, widespread labor market power is a dominant part of distortions in US manufacturing. This result helps explain the sharp fall in the manufacturing labor share, among other secular trends. Second, labor market power comes from technology, whereby plants capture most of the surplus of increased labor productivity. This result suggests scope for policy action to counteract technology-driven increases in labor market power.

Many questions remain about the nature of labor market power. First, we suggest future work to explore other Census data on non-manufacturing establishments in manufacturing firms to better understand how firm boundaries affect market power measurement. Second, there is considerable space remains to develop and estimate new models of production functions that better capture

trends in output elasticities, especially for periods of dramatic technological change. Third, our evidence on technological change is associative; researchers might better understand the causal mechanisms behind technological change using innovation surveys or quasi-experimental variation. Finally, we hope to quantify how these trends affect welfare, a challenge that involves careful modeling of plant and worker dynamics.

# References


Acemoglu, D. and P. Restrepo (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives 33*(2).

Ackerberg, D., K. Caves, and G. Frazer (2015). Identification properties of recent production function estimators. *Econometrica 83*(6).

Atalay, E., A. Hortaçsu, and C. Syverson (2014). Vertical integration and input flows. *American Economic Review 104*(4).

Autor, D. and D. Dorn (2013). The growth of low-skill service jobs and the polarization of the US labor market. *American Economic Review 103*(5).

Autor, D., D. Dorn, and G. Hanson (2013). The China syndrome: Local labor market effects of import competition in the United States. *American Economic Review 103*(6).

Autor, D., D. Dorn, L. Katz, C. Patterson, and J. Van Reenen (2020). The fall of the labor share and the rise of superstar firms. *Quarterly Journal of Economics 135*(2).

Autor, D. and A. Salomons (2018). Is automation labor-displacing? Productivity growth, employment, and the labor share. *Brookings Papers on Economic Activity*.

Azar, J., I. Marinescu, and M. Steinbaum (2022). Labor market concentration. *Journal of Human Resources 57*(S).

Azar, J. A., S. T. Berry, and I. Marinescu (2022). Estimating labor market power. Working paper, National Bureau of Economic Research.

Baker, J. and T. Bresnahan (1988). Estimating the residual demand curve facing a single firm. *International Journal of Industrial Organization 6*(3).

Basu, S. (2019). Are price-cost markups rising in the United States? A discussion of the evidence. *Journal of Economic Perspectives 33*(3).

Baumol, W. (1982). Contestable markets: An uprising in the theory of industry structure. *American Economic Review 72*(1).

Benmelech, E., N. Bergman, and H. Kim (2022). Strong employers and weak employees: How does employer concentration affect wages? *Journal of Human Resources 57*(S).

Berger, D., K. Herkenhoff, and S. Mongey (2022). Labor market power. *American Economic Review 112*(4).

Berry, S., M. Gaynor, and F. Scott Morton (2019). Do increasing markups matter? Lessons from empirical industrial organization. *Journal of Economic Perspectives 33*(3).

Bivens, J. and L. Mishel (2015). Understanding the historic divergence between productivity and a typical worker's pay: Why it matters and why it's real. *Economic Policy Institute*.

Blatter, M., S. Muehlemann, and S. Schenker (2012). The costs of hiring skilled workers. *European Economic Review 56*(1).

Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica 77*(3).

Bond, S., A. Hashemi, G. Kaplan, and P. Żoch (2021). Some unpleasant markup arithmetic: Production function elasticities and their estimation from production data. *Journal of Monetary Economics 33*.

Brill, M., C. Holman, C. Morris, R. Raichoudhary, and N. Yosif (2017). Understanding the labor productivity and compensation gap. *Beyond the Numbers 6*(6).

Byrne, D., J. Fernald, and M. Reinsdorf (2016). Does the United States have a productivity slowdown or a measurement problem? *Brookings Papers on Economic Activity*.

Charles, K. K., E. Hurst, and M. Schwartz (2019). The transformation of manufacturing and the decline in US employment. *NBER Macroeconomics Annual 33*(1).

De Loecker, J., J. Eeckhout, and G. Unger (2020). The rise of market power and the macroeconomic implications. *Quarterly Journal of Economics 135*(2).

De Loecker, J. and F. Warzynski (2012). Markups and firm-level export status. *American Economic Review 102*(6).

Dobbelaere, S. and J. Mairesse (2013). Panel data estimates of the production function and product and labor market imperfections. *Journal of Applied Econometrics 28*(1).

Edmond, C., V. Midrigan, and D. Y. Xu (2023). How costly are markups? *Journal of Political Economy 131*(7).

Elejalde-Ruiz, A. (2016). Manufacturing's big challenge: Finding skilled and interested workers. *Chicago Tribune*.

Elsby, M., B. Hobijn, and A. Şahin (2013). The decline of the US labor share. *Brookings Papers on Economic Activity*.

Flynn, Z., J. Traina, and A. Gandhi (2019). Measuring markups with production data. Working paper, Social Science Research Network.

Fort, T., J. Pierce, and P. Schott (2018). New perspectives on the decline of US manufacturing employment. *Journal of Economic Perspectives 32*(2).

Foster, L., J. Haltiwanger, and C. Syverson (2008). Reallocation, firm turnover, and efficiency: Selection on productivity or profitability? *American Economic Review 98*(1).

Gandhi, A., S. Navarro, and D. Rivers (2020). On the identification of gross output production functions. *Journal of Political Economy 128*(8).

Goos, M., A. Manning, and A. Salomons (2014). Explaining job polarization: Routine-biased technological change and offshoring. *American Economic Review 104*(8).

Hall, R. (1988). The relation between price and marginal cost in US industry. *Journal of Political Economy 96*(5).

Hall, R. (2018). New evidence on the markup of prices over marginal costs and the role of mega-firms in the US economy. Working paper, National Bureau of Economic Research.

Hashemi, A., I. Kirov, and J. Traina (2022). The production approach to markup estimation often measures input distortions. *Economics Letters 217*.

Hershbein, B., C. Macaluso, and C. Yeh (2019). Concentration in US local labor markets: Evidence from vacancy and employment data. Mimeo.

Karabarbounis, L. and B. Neiman (2014). The global decline of the labor share. *Quarterly Journal of Economics 129*(1).

Kehrig, M. and N. Vincent (2021). The micro-level anatomy of the labor share decline. *Quarterly Journal of Economics 136*(2).

Kirov, I., P. Mengano, and J. Traina (2023). Measuring markups with revenue data. Working paper, Social Science Research Network.

Klette, T. J. and Z. Griliches (1996). The inconsistency of common scale estimators when output prices are unobserved and endogenous. *Journal of Applied Econometrics 11*(4).

Lamadon, T., M. Mogstad, and B. Setzler (2022). Imperfect competition, compensating differentials, and rent sharing in the US labor market. *American Economic Review 112*(1).

Leduc, S. and Z. Liu (2023). Automation, bargaining power, and labor market fluctuations. Working paper, Federal Reserve Bank of San Francisco.

Levinsohn, J. and A. Petrin (2003). Estimating production functions using inputs to control for unobservables. *Review of Economic Studies 70*(2).

Manning, A. (2006). A generalised model of monopsony. *Economic Journal 116*(508).

Morlacco, M. (2020). Market power in input markets: Theory and evidence from French manufacturing. Working paper, MacMillan Center.

Olley, S. and A. Pakes (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica 64*(6).

Rinz, K. (2022). Labor market concentration, earnings, and inequality. *Journal of Human Resources 57*(S).

Robinson, J. (1933). *The Economics of Imperfect Competition*. Macmillan.

Shapiro, C. (2019). Protecting competition in the American economy: Merger control, tech titans, labor markets. *Journal of Economic Perspectives 33*(3).

Stansbury, A. and L. Summers (2017). Productivity and pay: Is the link broken? Working paper, National Bureau of Economic Research.

Stansbury, A. and L. Summers (2020). The declining worker power hypothesis: An explanation for the recent evolution of the American economy. *Brookings Papers on Economic Activity*.

Sussman, A. L. (2016). As skill requirements increase, more manufacturing jobs go unfilled. *Wall Street Journal*.

Syverson, C. (2011). What determines productivity? *Journal of Economic literature 49*(2).

Syverson, C. (2019). Macroeconomics and market power: Context, implications, and open questions. *Journal of Economic Perspectives 33*(3).

Traina, J. (2018). Is aggregate market power increasing? Production trends using financial statements. Working paper, Stigler Center.

# A Manufacturing Microdata

We use US Census Bureau production information from the Annual Survey of Manufactures and Census of Manufactures. The data include annual information on output, intermediates, labor, capital, and other characteristics at the plant level. We specifically use the Total Factor Productivity Beta Version 2 dataset, which applies basic cleaning procedures to the raw data and spans 1972 to 2014. This dataset underlies commonly disclosed productivity statistics, and is especially useful as a gold standard for replication. Our further cleaning procedure is as follows.

First, we select our sample following two principles: (1) We want to avoid inference driven by extreme outliers, and (2) we need available and consistent information on log output and inputs to eventually apply our panel estimators. For (1), we drop observations that have particularly large or small $\frac{R}{CM}$ or $\frac{CM}{WL}$ ratios, above the 99th or below the 1st percentile, calculated separately for each year. As these ratios directly enter our wedge calculations, this condition trims $\mathcal{M}$ and $\Delta$ outliers preemptively; we also surmise that outlier plants are likely to

have significant measurement error or radically different technologies. For (2), we limit our sample to observations with over 100 employees, as plants above this threshold are sampled each year with certainty. We also require positive output and inputs. After applying these filters to plant-year observations, if a plant has a gap in its panel (so that it has missing information between years of nonmissing information), we drop it entirely from our sample.

Second, we use external aggregate data to account for unmeasured productive inputs. In particular, plant surveys may not account for all nonproduction labor, such as managers at headquarter establishments. This problem may have grown over time as economies of scale increased. To address this concern, we merge in data from US KLEMS. These data combine information from the Input-Output Accounts produced by the Bureau of Economic Analysis (BEA) and the Bureau of Labor Statistics to derive harmonized measures of outputs and inputs for 65 subsectors, 19 of which are in manufacturing. We use these measures to proportionally scale our microdata output and input prices and quantities to match aggregate industry averages.

We measure output using the total value of shipments less sales from inventories: that is, output produced rather than output sold. We define intermediates as energy plus materials. For quantity measures, we deflate energy and materials with price indices from the NBER-CES Manufacturing Industry Database.

We measure production labor inputs as total production worker hours, and nonproduction labor inputs as total nonproduction workers employed (we do not have hours information for nonproduction workers). The average wage for nonproduction workers in our sample is over 40 USD per hour, about four times that of production workers. We therefore interpret $N$ as managerial and professional workers. To construct wages, we divide each plant's total production worker expenditures by its total production worker hours.

We construct capital as the sum of equipment and structures using the perpetual inventory method. We initialize each plant's capital stock to the book

value in the first Economic Census year in which it appears, then iterate forwards and backwards using investment data. We adjust capital stocks by their industry capital-utilization rates from the Federal Reserve's Industrial Production and Capacity Utilization survey.