Understanding High-Wage Firms:

Monopoly, Monopsony, and Bargaining Power^{*}

Horng Chern Wong⁺

July 2025

Abstract

I study how firm market power and worker bargaining power shape wages and welfare. Using French micro-data, I document patterns linking wages and firm market power that existing models cannot explain. A model in which firms produce vertically differentiated goods and share profits with workers explains those patterns. The model (a) reveals new challenges in estimating monopsony and bargaining power, proposing an alternative approach; (b) shows that the passthrough of firm-specific shocks to wages depends on the type of shock; (c) explains how markups shape firm wage premia; and (d) formalizes how strengthening worker bargaining power affects wages and welfare.

Keywords: wage inequality, firm heterogeneity, monopoly, monopsony, bargaining *JEL codes*: E24, J31, J42, L11

[†]Stockholm University, Department of Economics. Email: horngchernwong@gmail.com

^{*}First version: August 2019. Previously circulated as "Understanding High-Wage Firms". Special thanks to my dissertation advisors—Thijs van Rens, Roland Rathelot, and Dennis Novy—for their invaluable guidance. I am also very grateful to the editor, Katja Seim, as well as Chinhui Juhn and the anonymous reviewers, along with Jan Eeckhout, Lawrence Katz, Paul Klein, Francis Kramarz, Jay Euijung Lee, Attila Lindner, Isabelle Mejean, Aviv Nevo, Andrew Oswald, Vincent Sterk, and John Van Reenen, for their excellent comments and suggestions. For insightful conversations, I thank Jesper Bagger, Dan Bernhardt, Christine Braun, Danilo Cascaldi-Garcia, Dita Eckardt, Basile Grassi, Daniel Habermacher, Alice Kuegler, Thomas Le Barbanchon, Tim Sang Yoon Lee, Timo Leidecker, Norman Loayza, Roberto Pancrazi, Ayush Pant, Carlo Perroni, Josep Pijoan-Mas, Camilla Roncoroni, Nina Roussille, Alessandro Ruggieri, Marta Santamaria, Jesse Shapiro, Ben Smith, Douglas Staiger, Anna Stansbury, Liliana Varela, Marija Vukotic, Jonathan Yeo, Anders Åkerman, and the many seminar and conference participants. I gratefully acknowledge financial support from the ESRC, the ERC TRADENET grant (714597), and the Tore Browaldh Foundation, as well as data access from DARES, DGFiP, and INSEE.

1 Introduction

The concentration of economic activity among highly profitable firms has garnered increasing attention due to their impact on labor markets (Autor, Dorn, Katz, Patterson, and Van Reenen, 2020). These firms contribute to wage inequality by paying a wage premium for similar workers compared to less profitable firms (Card, Cardoso, Heining, and Kline, 2018). One line of research quantifies firms' labor market power as a determinant of profits and wages (Berger, Herkenhoff, and Mongey, 2022; Lamadon, Mogstad, and Setzler, 2022), but generally abstracts from product market power and worker bargaining power as alternative drivers. Another line quantifies product market power as a determinant of profits and Unger, 2020; Edmond, Midrigan, and Xu, 2023), but largely abstracts from imperfect labor market competition and bargaining, leaving firm wage premia and the welfare implications of bargaining power unaddressed.

This paper presents a structural framework where firms have both product and labor market power, and workers have collective bargaining power. A central theoretical insight is that the wedge between wages and the marginal revenue product of labor—the *labor wedge*—can be decomposed into three components: price-cost markups, monopsony markdowns, and bargaining power. The model also highlights identification challenges in distinguishing bargaining power from markdowns using standard approaches, and proposes a novel method to address them. In a quantitative version of the model, I show how firm heterogeneity—productivity, product quality, and non-wage amenities—and shocks to them affect firm wage premia. I further formalize the conditions under which strengthening worker power improves welfare. Using detailed French administrative data, which include measures of output prices and hours worked, I estimate markups, markdowns, bargaining power, and firm heterogeneity, and quantify their effects on wages and welfare.

My main findings highlight the importance of firms' product market characteristics and workers' bargaining power in determining firm wage premia and welfare. Two key empirical facts in the French data are inconsistent with a standard monopsony model—where differentially productive firms operate in competitive product markets and do not bargain with workers: (i) high-wage firms charge higher output prices and markups, and (ii) they pay workers a larger share of their marginal revenue product. My model accounts for these facts and yields several novel insights.

First, the elasticity of firm wage premia to idiosyncratic shocks depends on the nature of the shock. When products are more differentiated, consumers value quality over price, leading to weaker passthrough of productivity gains—but stronger passthrough of quality improvements—to wages. Second, although higher markups can amplify wage dispersion through rent-sharing, they also restrict output and labor demand, which dampens wage premia dispersion. Quantitatively, this dampening effect dominates. Third, product market power accounts for the majority—over 80%—of the welfare losses associated with firm market power. Finally, increasing worker bargaining power improves welfare mainly by offsetting the distortions induced by markups, though it increases wage inequality between firms. Giving workers full bargaining power closes less than one-third of the welfare gap between the decentralized market economy and the welfare-maximizing planner's economy.

In Section 2, I develop a flexible partial equilibrium framework with imperfect competition in both product and labor markets. Firms differ in productivity, product quality, and amenities. They face upward-sloping labor supply and downward-sloping product demand, but the model does not commit to a specific micro-foundation at this point.¹ Firms set prices as a markup over marginal cost. Wages are determined through bargaining between workers and firms, taking into account profits generated by imperfect competition. Workers at each firm bargain over wages *collectively* with their employer—in line with a key feature of French wage-setting institutions and a source of tractability for the model.² In equilibrium, a firm's wage equals its marginal revenue product of labor multiplied by a labor wedge. A higher labor wedge implies that workers capture a larger share of their marginal revenue product.

The first main contribution of this paper is to show that the *labor wedge* is driven by price markups, bargaining power, and monopsony markdowns. Prior work that quantifies firms' labor market power largely abstracts from markups and bargaining (Berger et al., 2022; Lamadon et al., 2022). My model nests monopsony as a special case—when workers have no bargaining power, the labor wedge corresponds to the

¹These may arise for reasons such as search frictions or product/workplace differentiation (Burdett and Mortensen, 1998; Gourio and Rudanko, 2014; Card et al., 2018)

²Bhuller, Moene, Mogstad, and Vestad (2022) show that firm-level collective bargaining is prevalent among OECD countries. The model does not feature individual bargaining and I discuss this limitation in Appendix C.1.

monopsony markdown. However, when workers have some bargaining power and firms have product market power, workers can capture a share of the profits from markups through wages. This *rent-sharing* channel implies that the labor wedge increases with markups. Recent studies that integrate product and labor market power focus on wage-posting models and do not feature bargaining, thus abstracting from this rent-sharing channel (Kroft, Luo, Mogstad, and Setzler, 2024; Deb, Eeckhout, Patel, and Warren, 2024). In those models, markups only reduce labor demand as firms restrict output, thereby reducing wages by moving down the labor supply curve.

One key implication is that workers' collective bargaining power can offset the effects of firm product and labor market power on labor demand. Without bargaining, firms exploit monopsony power by restricting employment and product market power by limiting output. Worker power shifts some profits to workers, compelling firms to expand production and employment. I show that sufficiently strong worker power can eliminate the labor wedge entirely. Recent work quantifying the macroe-conomic impact of markups and markdowns does not incorporate bargaining (Berger et al., 2022; Edmond et al., 2023). Strengthening worker power may attenuate firm market power and improve welfare. I explore this idea quantitatively in Section 7.

In Section 3, I show that my model raises new challenges for separately estimating monopsony markdowns and bargaining power. Recent advances in the literature show that, using the "production approach", monopsony markdowns can be directly estimated as the wedge between wages and marginal revenue products of labor (Yeh, Macaluso, and Hershbein, 2022). However, my model clarifies that this approach only recovers markdowns when workers have no bargaining power. Meanwhile, the "rentsharing" literature measures worker bargaining power by estimating the passthrough of profits to wages, using instruments such as export demand or innovation-based shocks (Van Reenen, 1996; Kline, Petkova, Williams, and Zidar, 2019; Garin and Silverio, 2024). This approach often relies on wage bargaining models that do not feature firm monopsony power (Card et al., 2018).³ In such models, reservation wages are independent of firm-specific shocks, ensuring the validity of instruments for profits. In contrast, in my model, workers without bargaining power earn endogenous, firmspecific *monopsony wages* that move with the same shocks affecting firm profits. As a

³One exception is Kline et al. (2019), who motivate their analysis with a monopsony model. Card et al. (2018) also explain how a monopsony model could interpret those rent-sharing estimates.

result, existing instruments violate the exclusion restriction in my setting.

The model suggests a new path forward to disentangling workers' collective bargaining power from monopsony markdowns. Rather than estimating bargaining power from a wage equation as in the rent-sharing literature, I exploit the labor wedge equation, which avoids the problem of unobserved monopsony wages. This approach relies on the property that labor wedges increase with markups when workers have bargaining power. Given measures of labor wedges and markups, the main identification challenge is addressing unobserved monopsony markdowns, which are firmspecific and endogenous to firm size and amenities. While firm size is observed in the data, amenities are not. To address this, I propose a theory-consistent control function approach: firms' wage bill and employment jointly proxy for amenities, based on the idea that firms with better amenities attract more workers at a given wage. This monotonicity assumption underpins my control function strategy. With an estimate of bargaining power, firm-specific monopsony markdowns can then be recovered.

Implementing my approach begins with measuring labor wedges and markups. I do this by applying the production approach of De Loecker and Warzynski (2012) and Yeh et al. (2022). This approach involves estimating production functions using a control function method to account for unobserved firm productivity, following Ackerberg, Frazer, and Caves (2015). I then separate labor wedges from markups by noting that labor wedges specifically distort labor demand, while markups distort all input demands. This approach offers some flexibility as it does not require researchers to commit to one specific market structure across product and labor markets.

Section 4 presents the detailed French micro-data I use to estimate markups, markdowns, bargaining power, and firm wage premia. The firm balance sheet panel data contains essential variables for production function estimation. To measure labor inputs and to estimate firm wage premia, I use the matched employer-employee panel data, which has the advantage of including data on hours worked, allowing me to account for differences in hours across firms. I complement these datasets with survey data on firm×product-level prices for manufacturing firms. The survey data helps address the common challenge of unobserved input and output prices in production function estimation (De Loecker, Goldberg, Khandelwal, and Pavcnik, 2016).

Section 5 presents my main empirical findings. I estimate a median labor wedge of

0.62 and a median monopsony markdown of 0.45 for French manufacturing firms. The difference is due to worker bargaining power and markups. I estimate that workers capture about 12% of firm profits, similar to the typical range summarized in Jäger, Schoefer, Young, and Zweimüller (2020). These results suggest that not accounting for worker bargaining power would *understate* firm monopsony power.

Comparing my labor wedge estimates with those from existing studies that use the production approach underscores the importance of accounting for hours worked in measures of labor inputs. My estimates are lower than those of Yeh et al. (2022) and Mertens (2020), who measure labor inputs by employment and report labor wedges of around 0.70 for median US and German manufacturing firms. When I measure labor inputs as employment, I find median labor wedges of 0.71, significantly narrowing the gap between my original estimates and those from existing literature. This discrepancy reflects an upward bias in estimated labor output elasticities when hours are omitted, as employment understates total labor input and larger firms tend to have longer average work hours.

My estimates reveal two novel empirical relationships between firm wage premia and firm market power. First, high-wage firms tend to charge higher output prices and markups, a pattern that persists even within narrow industries and cannot be explained by models with only productivity heterogeneity. Second, high-wage firms tend to have higher labor wedges, paying a larger share of marginal revenue products as wages. This finding is inconsistent with monopsony models without bargaining or product market power (Card et al., 2018), where labor wedges equal monopsony markdowns and typically decline with firm wage premia. My model accounts for this pattern through positive worker bargaining power, which allows workers to capture part of the product market rents.

In Section 6, I develop a quantitative general equilibrium version of my model to quantify the role of firm market power, worker bargaining power, and the sources of firm heterogeneity in shaping wages and welfare. In this model, product markets are oligopolistic and labor markets oligopsonistic, following Atkeson and Burstein (2008) and Berger et al. (2022). The model implies that markups and markdowns are determined by firms' product and labor market shares. I use the empirical counterparts of these structural relationships to calibrate the underlying preference parameters.

A key advantage of the French data is that it includes output prices and hours worked, which allows me to measure both product quality and amenities. I measure these based on the model-implied relationship between firms' sales and wage-bill market shares and their corresponding output prices and (hourly) firm wage premium.

The model's first contribution is to show how firm heterogeneity and firm-specific shocks shape wage premia. Section 7 decomposes the passthrough of a shock into three parts: (i) a *direct effect* on labor demand; (ii) a *market power effect* from changes in markups and markdowns, governed by worker bargaining power; and (iii) a *firm size effect* capturing general equilibrium spillovers. A key insight is that the size of the direct effect depends on the *nature* of the shock. Productivity shocks generate larger passthrough to wages than equally sized quality shocks when products are close substitutes, because sales are more sensitive to prices than to quality.

Quantitatively, a 1% productivity shock has a passthrough of 0.99, compared to 0.24 for a quality shock, mainly reflecting high substitutability across varieties. Nonetheless, quality heterogeneity generates more wage variance than productivity heterogeneity. Most empirical studies, which do not distinguish between shocks, find pass-through below 0.3 (Card et al., 2018; Lamadon et al., 2022). My findings suggest that the wage effects of innovation depend on whether innovations improve productivity or product quality, and on the degree of product differentiation.

The model's second contribution is to formalize and quantify how workers' collective bargaining power affects welfare when firms have product and labor market power—an idea long discussed by economists but rarely formalized (Stansbury and Summers, 2020). Comparing the decentralized and welfare-maximizing social planner's equilibra, I show that markups and markdowns reduce welfare by (i) acting as uniform taxes on aggregate labor demand, and (ii) generating misallocation. Raising worker bargaining power helps correct both distortions by redistributing profits to workers, which counteracts firm market power and compels firms to increase production and employment. However, even full worker power cannot fully restore efficiency, since markups also distort non-labor input demand.

Section 7 shows that markups and markdowns reduce welfare by 46% in consumptionequivalent terms, with markups alone accounting for over 80% of the loss. Misallocation explains about 63% of the welfare cost of markups. Markups also affect wages, potentially raising them through rent-sharing but may also reduce them by restricting labor demand. Quantitatively, the latter dominates: equalizing markups raises both average wages and wage variance by 39% and 97%, while increasing welfare by 24%.

Strengthening worker bargaining power improves welfare, but less so than eliminating markups entirely. Increasing bargaining power from 12% to 50% yields wage gains comparable to those from removing markups, but achieves only one-quarter of the welfare gain (10% vs. 38%). Granting workers full bargaining power raises welfare by 13%, closing less than a third of the gap to the planner's equilibrium. The remainder reflects distortions to non-labor input demand that bargaining alone cannot address.

2 A Partial Equilibrium Model of Firm & Worker Power

I now set up a model of the labor market in which heterogeneous firms have labor and product market power, and workers have bargaining power. The model serves several purposes: (i) to structurally interpret regression-based estimates of firm wage premia; (ii) to clarify the role of worker bargaining power in determining labor demand and wages; (iii) to guide how firm market power and worker bargaining power may be estimated from the data (in Section 3). In Section 6, I close the model in general equilibrium and impose further structure to quantify the implications of firm and worker power for wages and welfare.

2.1 Model environment

Labor supply. Let Φ_j be the piece-rate wage per efficiency unit of labor paid by firm *j*. A worker *i* with efficiency E_i obtains a wage $W_{ij} = E_i \Phi_j$. Taking logs, this wage equation has a log-additive structure reminiscent of the "AKM" regression due to Abowd, Kramarz, and Margolis (1999): $w_{ij} = e_i + \phi_j$, where lowercase letters denote variables in logs. The piece-rate wage (ϕ) is the *firm-specific wage (premium)*.

Efficiency units of labor in the firm are $H_j = \bar{E}_j L_j$, where \bar{E}_j denotes average efficiency and L_j denotes amount of labor. Let the upward-sloping labor supply curve facing each firm be $H_j = \mathcal{H}(\Phi_j, A_j)$. The firm-specific labor supply shifter A_j represents non-wage amenities and is taken as given by the firm. I assume that the labor supply function is twice differentiable and monotonically increasing in Φ_j and A_j .

Product demand. Let the downward-sloping demand curve for firm j's output be $Y_j = \mathcal{G}(P_j, D_j)$. The price charged by firm j is P_j . Firms take as given the demand shifter for its goods D_j , capturing its product quality. The goods demand function is twice differentiable and monotonically decreasing in P_j and increasing in D_j .

Production technology. Firms operate a general production function with diminishing marginal returns to each input $Y_j = \Omega_j F(K_j, M_j, H_j)$. Ω_j is the Hicks-neutral productivity, K_j are physical capital, M_j are material inputs, and H_j are units of effective labor. The production function is twice differentiable and satisfies $F(K_j, M_j, 0) = 0$. The markets for capital and materials are perfectly competitive with prices P_k and P_m .

Wage determination. An important institutional feature of wage bargaining in France is that firms with at least 50 employees are legally required to bargain annually with their employees, who are represented by labor union representatives (see Appendix A). That is, workers bargain *collectively* with their employer *j* on a per-period basis. Firm-level collective bargaining is not unique to France—as Bhuller et al. (2022) show, this institutional feature has become increasingly common among OECD countries.

Modeling firm-level collective bargaining between workers and their employers requires defining the workers' collective threat if wage negotiations were to break down. I assume that workers can threaten to go on a collective one-period strike, halting production and leading to zero profits for firms at the cost of zero wages for that period. Since the model is static and bargaining occurs every period, the bargaining problem with the same workers is repeated next period. As we will see, this collective threat allows workers to demand a wage above the level that would prevail in a monopsony, with the extent depending on workers' bargaining strength.⁴ This approach is similar to recent work on labor unions. For example, Taschereau-Dumouchel (2020) models unionized firms as paying higher wages than non-unionized firms because of workers' ability to collectively quit into unemployment if no agreement is reached.

I assume that bargaining is efficient in the sense that workers and firms jointly decide on wages, prices, materials, and capital to maximize total profits, taking into account the product demand curve and labor supply curve.⁵ Let Π_i be firm profits.

⁴Another reason workers may be able to demand wages above the monopsony level is if workers' *individual* outside options improve. I discuss this point further in Appendix C.1.

⁵When capital is a production input and capital investments are sunk, an important concern is whether workers can holdup their employers and extract rents. However, recent evidence suggests

Worker bargaining power is $\kappa \in [0, 1]$. The collective bargaining problem is:

$$\max_{\Phi_j, P_j, M_j, K_j} \left(\Phi_j H_j \right)^{\kappa} \left(\Pi_j \right)^{1-}$$

subject to $H_j = \mathcal{H}(\Phi_j, A_j)$, $Y_j = \mathcal{G}(P_j, D_j)$, and $Y_j = \Omega_j F(K_j, M_j, H_j)$. The firm's profit is $\Pi_j = P_j Y_j - \Phi_j H_j - P_m M_j - P_k K_j$. When workers have no bargaining power ($\kappa = 0$), they have no influence on the firm's input and price/wage-setting decisions.

2.2 How bargaining affects firm wage premia and labor demand

2.2.1 Firm wage premia

Solving the bargaining problem gives the following firm-specific wage (premium):

$$\Phi_{j} = \kappa \left(\underbrace{\frac{P_{j}Y_{j} - P_{m}M_{j} - P_{k}K_{j}}{H_{j}}}_{\text{Profits per effective labor}} \right) + (1 - \kappa) \underbrace{\lambda_{j}MRPH_{j}}_{\text{Monopsony}}$$
(1)

The first term represents *profits per effective labor*, also frequently referred to as "quasirents" in the rent-sharing literature (Card et al., 2018). The second term is the *monopsony wage*, which is a markdown (λ_j) below the marginal revenue product of effective labor (*MRPH_j*). The monopsony markdown $\lambda_j = \frac{\xi_j}{1+\xi_j}$ is a function of labor supply elasticities $\xi_j = \xi(H_j, A_j)$, whose specific functional form depends on the microfoundation for the labor supply curve.

Equation (1) shows that the firm wage premium is a weighted average of two common wage-setting mechanisms: a pure bargaining outcome and a pure monopsony outcome. When workers have no bargaining power ($\kappa = 0$), firms set wages unilaterally taking into account the labor supply curve. As a result, workers are paid the monopsony wage, as is the case in wage-posting models of the labor market (Berger et al., 2022; Lamadon et al., 2022; Kroft et al., 2024). When workers have some bargaining power ($\kappa \in (0, 1)$), they are able to extract a share of the firm profits. Therefore, bargained wages are a top-up over the monopsony wage. When workers have full bargaining power ($\kappa = 1$), they capture the full profits. Thus, higher worker bargaining power redistributes profits generated by imperfect product and labor market competition from firms to workers. This redistribution of profits also impacts firm

that such holdup problems tend to be small (see, for example, Card, Devicienti, and Maida (2014)).

labor demand, as I show next.

2.2.2 Labor demand

To see how markups, markdowns, and bargaining affect firm labor demand, consider a different representation of the firm wage premium equation (1):

$$\Phi_j = \Lambda_j M R P H_j = \Lambda_j \mu_j^{-1} M P H_j \tag{2}$$

where the marginal revenue product of effective labor is $MRPH_j = \mu_j^{-1}MPH_j = \mu_j^{-1}\alpha_{h,j}\frac{P_jY_j}{H_j}$. Let $MPH_j = \alpha_{h,j}\frac{P_jY_j}{H_j}$ denote the *value* of the marginal product of labor. $\frac{P_jY_j}{H_j}$ is the revenue per effective worker, $\alpha_{h,j}$ is the output elasticity with respect to labor inputs, and $\mu_j = \frac{\rho_j}{\rho_j - 1}$ is the *price-cost markup*, which is a function of the price elasticities of demand $\rho_j = \rho(Y_j, D_j)$. The *labor wedge* (Λ_j) is the share of $MRPH_j$ that workers receive as wages.

Markups as negative labor demand shifters. Notice that in equation (2) markups directly drive a wedge between wages and marginal products of labor (MPH_j) . Because markups enter inversely, a higher markup represents a negative labor demand shift, capturing the idea that monopolists *reduce output*, and hence labor demand. In a model where firms have both product and labor market power but workers have no bargaining power, all else equal, a markup-induced negative labor demand shift moves firms down the labor supply curve, thus reducing wages (Kroft et al., 2024).

The role of wage bargaining and firm market power in determining the labor wedge. A key insight of the model is that the difference between wages and the marginal revenue product of labor—the *labor wedge* (Λ_i)—can be expressed as:

$$\Lambda_j = \kappa \left(1 - \frac{\alpha_{m,j} + \alpha_{k,j}}{\mu_j} \right) \frac{\mu_j}{\alpha_{h,j}} + (1 - \kappa) \lambda_j$$
(3)

The first term reflects the profits per unit of MRPH ("product market rents" henceforth) captured by workers through bargaining—a *rent-sharing* term. In contrast with the role of markups as negative labor demand shifters, this term increases with markups, as higher markups increase the profits that workers can appropriate. Thus, markups can also act as positive labor demand shifters by raising the labor wedge (Λ_j). This rent-sharing term introduces a new channel through which markups affect labor demand and wages relative to models without bargaining (Kroft et al., 2024).

The second term arises from the monopsony markdown (λ_j) , representing firm monopsony power when workers have no bargaining power ($\kappa = 0$). When $\kappa = 0$, the labor wedge is equal to the monopsony markdown $\Lambda_j = \lambda_j$, as workers do not capture any profits. When $\kappa > 0$, the labor wedge exceeds the monopsony markdown. Thus, the labor wedge captures the idea that worker bargaining power can offset firm monopsony power, with pure monopsony as a special case (Lamadon et al., 2022). Additionally, since the labor wedge increases with markups, higher markups make worker bargaining power more effective at countering firm monopsony power.

3 Estimating Markups, Markdowns, Bargaining Power

The model in Section 2 relates firm wage premia (Φ_j) to labor wedges (Λ_j), markups (μ_j), markdowns (λ_j), and worker bargaining power (κ). This section provides a detailed description of my approach to estimating these objects.

3.1 Estimating firm wage premia

A common method for estimating firm wage premia is to recover firm effects from an AKM regression (Abowd et al., 1999), typically assuming firm effects are *fixed over time* and identifying them through worker mobility across firms. However, as I explain in the next subsection, measuring firms' efficiency units of labor over time requires estimating *time-varying* firm wage premia. A practical challenge is the limited worker mobility in short panels, which leads to noisy firm effect estimates and an upward bias in their variance. To address both the need for time variation and the lack of mobility, I implement the k-means classification approach of Bonhomme, Lamadon, and Manresa (2019) ("BLM" henceforth).⁶

I first classify firms into groups using a k-means clustering algorithm, then estimate a version of the AKM regression replacing firm effects with firm-group effects. Specifically, I estimate the following regression:

$$\ln W_{it} = \chi'_{it}\beta + \iota_i + \phi_{g(j(i,t),t)t} + \nu_{it}$$

⁶I compare the estimated variance of firm wage premia between the Bonhomme et al. (2019) k-means clustering and Kline, Saggio, and Solvsten (2020) leave-out approaches in Table 9.

where *i* denotes the individual, j(i, t) denotes the firm that employs *i* at time *t*, g(j, t) denotes the group of firm j at time *t*, ι_i are worker fixed effects, $\phi_{g(j(i,t),t)}$ are firm-group effects that vary by *t*, and χ_{it} includes age polynomials and part-time status. When there are as many firm-groups as there are firms, this regression converges to the AKM regression. The firm-group fixed effects are identified by workers who switch between firm-groups. Relative to the AKM regression, this procedure has the advantage that it substantially increases the number of switchers used to identify firm-group effects, enabling wage premia to be more precisely estimated.⁷

To classify firms with similar wage premia, I group firms based on the similarity of their internal wage distributions. Conditional on a log-additive wage structure, firms with similar firm and worker effects should exhibit similar internal wage distributions. If two firms have similar distributional shapes but different average wages, they likely differ in firm effects. Conversely, if average wages are similar but distributional shapes differ, they are assigned to different groups. In practice, I apply the clustering algorithm by 2-digit sector over overlapping 2-year windows, allowing wage premia to vary over time. More details on the clustering procedure and the underlying restrictions of the AKM framework are provided in Appendix B.1.

3.2 Estimating labor wedges and price-cost markups

Estimation approach in theory. Having estimated firm wage premia, I now apply the production approach to separately estimate markups and labor wedges (De Loecker and Warzynski, 2012; Yeh et al., 2022). I start by estimating production functions to obtain output elasticities. Let the production function for each sector *s* be $y_{jt} = f_s(k_{jt}, m_{jt}, h_{jt}; \beta) + \omega_{jt}$. Lowercase letters represent the natural log counterparts of variables written in uppercase letters, β is the set of production function parameters, and ω_{jt} is the firm's Hicks-neutral productivity. The firm-specific output elasticities with respect to capital, material inputs, and effective labor are: $\alpha_{k,jt} := \frac{\partial y_{jt}}{\partial k_{jt}}, \alpha_{m,jt} := \frac{\partial y_{jt}}{\partial m_{jt}}$, and $\alpha_{h,jt} := \frac{\partial y_{jt}}{\partial h_{jt}}$. Consistent with the model, I measure effective labor as $H_{jt} = \bar{E}_{jt}L_{jt}$, where the average efficiency of workers per hour is the ratio of the firm's average wage to the firm wage premium at a given time t, $\bar{E}_{jt} = \frac{\bar{W}_{jt}}{\Phi_{ij}}$, and $\log \bar{E}_{jt}$ is normalized to have

⁷See Bonhomme, Holzheu, Lamadon, Manresa, Mogstad, and Setzler (2023) for a systematic assessment of the importance of clustering firms before estimating firm effects.

a mean of 0 in the cross-section.^{8,9}

Next, I disentangle firms' markups from their labor wedges using the fact that markups distort the demand for all inputs, while labor wedges only distort labor demand. This requires: (i) choosing a flexible input, and (ii) assuming firms are price takers for that input. Assuming materials meet these conditions, markups are the sole distortion to material demand, and labor wedges are the sole distortion to labor demand *relative* to materials. Markups and labor wedges can then be expressed as:

$$\mu_{jt} = \alpha_{m,jt} \frac{P_{jt} Y_{jt}}{P_{m,t} M_{jt}} \quad \text{and} \quad \Lambda_{jt} = \frac{\Phi_{jt} H_{jt}}{P_{m,t} M_{jt}} \cdot \frac{\alpha_{m,jt}}{\alpha_{h,jt}} = \frac{W_{jt} L_{jt}}{P_{m,t} M_{jt}} \cdot \frac{\alpha_{m,jt}}{\alpha_{h,jt}}$$
(4)

Thus, both can be separately identified using data on the material cost share and the wage-bill-to-materials ratio, along with estimates of output elasticities. In Appendix B.7 and B.8, I examine the assumptions that material inputs are flexible and firms take their prices as given.

Identification challenges. Obtaining estimates of output elasticities requires estimating production functions, which entails several challenges: unobserved firm productivity and unobserved output/input prices. I now address them in turn.

First, firm productivity ω_{jt} is unobserved but determines input choices, implying that OLS estimates of production function parameters are biased. To address this, I implement the control function approach (Ackerberg et al., 2015), which corrects for this endogeneity by allowing researchers to indirectly observe firm productivity—through the inversion of firms' optimal demand for a fully flexible input.

Second, when estimating production functions, researchers often do not observe firms' output and input prices, which may lead to biased estimates of output elasticities (Klette and Griliches, 1996; De Loecker and Goldberg, 2014). To address unob-

⁸As I discuss in Appendix C.1, my model abstracts from the sequential auction wage-setting mechanism of Cahuc, Postel-Vinay, and Robin (2006), which allows incumbent workers and new hires—who are otherwise identical—to receive different wages as incumbents accumulate outside offers. In such cases, equation (2) may be misspecified. However, in Appendix B, I show that my empirical findings are robust to focusing on hiring wages only, following the approach of Di Addario, Kline, Saggio, and Solvsten (2020), thereby allowing for differential pay between incumbents and new hires.

⁹Implicit in the efficiency units specification of the production function, I assume that worker types are perfect substitutes, although average worker efficiency and firm productivity are complements. When workers are imperfect substitutes, the log-additive AKM regression is misspecified—an interaction term between the worker and firm effect needs to be present. I address this issue in two ways. First, I show that when workers change employers, the changes in AKM firm effects are very similar to the unconditional wage changes, following Sorkin (2018). Second, in Appendix E, I provide an extension that distinguishes between two skills. I find that the results of this extension are similar to those reported in Section 5.

served heterogeneity in output prices, I use data on output prices from a large survey of French manufacturing firms (see Section 4 for details). For unobserved heterogeneity in input prices—specifically capital and materials—I follow the approach of De Loecker et al. (2016), who show that, under the assumption that firms are pricetakers in input markets, output prices can be used to control for input price variation. In my model, however, firms have monopsony power in labor markets, leading to unobserved monopsony markdowns. To ensure the validity of the De Loecker et al. (2016) approach in this context, I additionally control for firm wage premia. Further details on how I address unobserved output and input prices are provided in Appendix B.3.

Input price control function. To aid the exposition that follows, I begin with a few notational definitions. Let $\tilde{\mathbf{x}}_{jt} = {\tilde{k}_{jt}, \tilde{m}_{jt}}$ denote capital and material input expenditures and $\mathbf{p}_{\mathbf{x},jt} = {p_{k,jt}, p_{m,jt}}$ the corresponding input prices. When firms are pricetakers in the markets for these inputs, a control function for their unobserved prices can be expressed as $B_s(\mathbf{p}_{\mathbf{x}}(p_{jt}, \mathbf{Z}_{jt}), \tilde{\mathbf{x}}_{jt}, h_{jt}; \beta, \zeta)$, following De Loecker et al. (2016). The vector β denotes production function parameters and ζ input price parameters. The vector \mathbf{Z}_{jt} contains firm wage premia and a set of (5-digit) sector×location fixed effects.

Productivity control function. The first step in deriving the productivity control function is to obtain the first-order condition for materials: $\tilde{m}_{jt} = m_s(\omega_{jt}, \tilde{k}_{jt}, h_{jt}, \mu_{jt}, p_{jt}, \mathbf{Z}_{jt})$. I invert this function to obtain the control function for productivity:¹⁰

$$\omega_{jt} = \omega_s(h_{jt}, k_{jt}, \tilde{m}_{jt}, \mu_{jt}, p_{jt}, \mathbf{Z}_{jt})$$
(5)

This relies on the following assumptions:

<u>Scalar unobservability</u>. The firm's idiosyncratic Hicks-neutral productivity ω_{jt} is the only unobservable variable determining firm material input demand.

Monotonicity. Conditional on the variables in the control function, material input demand is monotonically increasing in ω_{it} .

The productivity control function (5) highlights a key identification challenge in settings with imperfect competition: it depends on prices and markups, $\mu_{jt} = \mu(P_{jt}, D_{jt})$, which are typically unobserved. When prices and markups are not observed, the material input demand function generally cannot be inverted—violating the scalar un-

¹⁰Under the assumption that firms are pricetakers for capital and material inputs, De Loecker et al. (2016) show that these input price differences are absorbed by output prices and a set of sector-location fixed effects. Therefore, they do not appear directly in the productivity control function.

observable assumption—unless all variation in prices and markups is driven by ω_{jt} . In the French data, I observe output prices. Additionally, if markup heterogeneity is driven by ω_{jt} or by regional and sectoral differences in product market competition, these are absorbed by the control function.

However, variation in idiosyncratic demand D_{jt} , uncorrelated with TFP, can still induce markup variation. To address this, I impose more structure on the drivers of markups and include additional controls for them. Drawing on models of oligopolistic competition (Edmond, Midrigan, and Xu, 2015), I control for export status and market shares. Motivated by customer capital models (Gourio and Rudanko, 2014), I include firm age. Finally, I include a third-order polynomial in output prices to account for potential nonlinearities in the price-markup relationship. The identifying assumption is that these controls adequately capture markup variation unrelated to TFP.

Estimation. I estimate translog production functions for each 2-digit manufacturing sector, allowing output elasticities to vary across firms and over time:

$$y_{jt} = \beta_{k,s}k_{jt} + \beta_{h,s}h_{jt} + \beta_{m,s}m_{jt} + \beta_{kk,s}k_{jt}^2 + \beta_{hh,s}h_{jt}^2 + \beta_{mm,s}m_{jt}^2 + \beta_{kh,s}k_{jt}h_{jt} + \beta_{km,s}k_{jt}m_{jt} + \beta_{hm,s}h_{jt}m_{jt} + \beta_{khm,s}k_{jt}h_{jt}m_{jt} + \omega_{jt}$$

I maintain the assumption that production functions are the same across firms within 2-digit sectors and are constant over time. In Appendix B.6, I discuss the choice of estimating translog production functions in detail.

I estimate the production functions following the two-step GMM approach (Ackerberg et al., 2015). In the first-stage, I combine the productivity and input price control functions with the production function and estimate the following by OLS:

$$y_{jt} = \Psi_s(k_{jt}, h_{jt}, \tilde{m}_{jt}, p_{jt}, \mathbf{Z}_{jt}; \beta, \zeta) + \epsilon_{jt}$$
(6)

With a slight abuse of notation, the vector \mathbf{Z}_{jt} now also includes year effects, export status, market shares, and firm age, in addition to firm wage premia and sector×location fixed effects. This step estimates and removes the residual term ϵ_{jt} , capturing measurement error and productivity shocks that are unobserved by the firm and are therefore orthogonal to input choices.¹¹

In the second stage, I estimate the production function parameters β and input-

 $^{^{11}}$ I approximate $\Psi(.)$ with a third-order polynomial using each variable except dummy variables.

price-related parameters ζ by forming moment conditions. Firm productivity ω_{jt} can be written as a function of the parameters to be estimated { β , ζ }:

$$\omega_{jt}(\beta,\zeta) = \hat{\Psi}_{jt} - f_s(\tilde{k}_{jt}, h_{jt}, \tilde{m}_{jt}; \beta) - B_s(\mathbf{p}_{\mathbf{x}}(p_{jt}, \mathbf{Z}_{jt}), \tilde{\mathbf{x}}_{jt}, h_{jt}; \beta, \zeta)$$

Let the law of motion for the log of Hicks-neutral productivity be:

$$\omega_{jt} = g_s(\omega_{jt-1}) + \omega_{jt} \tag{7}$$

where $g_s(.)$ is a flexible function and ω_{jt} is a productivity shock.¹² Combining equation (6) and the law of motion for productivity (7) gives productivity shocks ω_{jt} as a function of the parameters of interest:

$$\omega_{jt}(\beta,\zeta) = \omega_{jt}(\beta,\zeta) - g_s(\omega_{jt-1}(\beta,\zeta))$$

I then form the following moment conditions:

$$E[\omega_{jt}(\beta,\zeta)\mathbf{X}_{jt}] = \mathbf{0}$$

where X_{jt} includes current and lagged capital, lagged effective labor, lagged material inputs, lags of the squared factor inputs, lagged interaction terms between the factor inputs, lagged output prices, lagged market shares, lagged export status, lagged firm age, lagged wage premia, and the interaction terms between lagged output prices and lagged factor inputs. Under the timing assumption on the productivity process, these lagged variables are orthogonal to current productivity shocks. In addition, capital inputs are assumed to be dynamic and pre-determined, therefore firms' current capital input demand are also orthogonal to current productivity shocks.

An important caveat to note is that identifying material output elasticities is challenging: Gandhi, Navarro, and Rivers (2020) show that, with perfectly competitive product markets, the only source of variation that identifies material output elasticities is time-series variation in material prices. I address this concern in Appendix B.9.

3.3 Estimating bargaining power and monopsony markdowns

Section 2 shows that labor wedges (Λ) are not generally equal to monopsony markdowns (λ), except in the absence of worker bargaining power (κ). Disentangling the

¹²I let $g_s(.)$ be linear. I approximate B(.) with a third-order polynomial using each variable except dummy variables. Dummy variables enter linearly.

two requires knowledge of κ . I now discuss the identification challenges commonly encountered when estimating κ , the additional challenge posed by my model, and a novel estimation approach that addresses this challenge.

Common identification challenges and existing solutions. The rent-sharing literature typically estimates the elasticity of wages with respect to measured profits per worker (Card et al., 2018), based on a wage equation structurally similar to equation (1). The key distinction is that, in the absence of worker bargaining power ($\kappa = 0$), workers receive a reservation wage $\Phi_j^r = \Phi^r(\Phi^c, A_j)$, which reflects a competitive outside option (Φ^c) and a firm-specific non-wage amenity (A_j). In contrast, in equation (1), if $\kappa = 0$, workers receive the monopsony wage $\Phi_j^{monop} = \lambda_j MRPH_j$.

For both wage equations, an OLS regression of firm-level average wages ($W_j = \bar{E}_j \Phi_j$) on quasi-rents ($Q_j^{rent} = \frac{P_j Y_j - P_k K_j - P_m M_j}{H_j}$) will give an upward-biased estimate of κ . The reasons are: (IC1.) unobserved worker heterogeneity, (IC2.) common shocks that shift firms' labor demand, and (IC3.) differences in amenities across firms.

To address IC1, the literature typically focuses on rent-sharing estimates for workers who do not switch firms ("stayers"), using matched worker-firm panel data to control for unobserved worker heterogeneity. To address IC2 and IC3, the literature proposes a set of instruments for Q_j^{rent} and makes the case that they are relevant, idiosyncratic (to address IC2), and orthogonal to firm amenities (to address IC3).¹³

A new identification challenge. Estimating κ using the more general wage equation (1) introduces an additional identification challenge: any idiosyncratic shock that shifts Q_j^{rent} will also affect the monopsony wage Φ_j^{monop} at the same time, implying that the exclusion restriction cannot be satisfied. This is because both Q_j^{rent} and Φ_j^{monop} are determined by firm productivity, product quality, and amenities. Thus, applying existing instruments to estimate worker bargaining power will yield biased estimates of κ in the context of my model.

An alternative approach to estimating worker bargaining power. I propose a theoryconsistent approach to estimating κ that avoids the presence of unobserved monopsony wages. Instead of using the wage equation (1), my approach exploits the structural relationship between price markups and labor wedges given by equation (3).

¹³See Card et al. (2018) for a survey of this literature. Recent work used idiosyncratic export demand shocks or patent shocks as instruments (Garin and Silverio, 2024; Kline et al., 2019).

Following equation (3), define product market rents as $\tilde{\mu}_{jt} \equiv \left(1 - \frac{\alpha_{m,j} + \alpha_{k,j}}{\mu_j}\right) \frac{\mu_j}{\alpha_{h,j}} = \left(\frac{P_{jt}Y_{jt} - P_{m,t}M_{jt} - P_{k,t}K_{jt}}{P_{m,t}M_{jt}}\right) \frac{\alpha_{m,jt}}{\alpha_{h,jt}}$. Suppose that labor wedges (Λ_{jt}), markdowns (λ_{jt}), and output elasticities ($\alpha_{h,jt}$, $\alpha_{m,jt}$) are observed.

Then, according to equation (3), conditional on monopsony markdowns (λ_{jt}), variation in product market rents ($\tilde{\mu}_{jt}$) identify worker bargaining power. The intuition is straightforward: when workers have bargaining power, they capture a share of rents generated by markups. When $\kappa = 0$, variation in $\tilde{\mu}_{jt}$ do not directly affect Λ_{jt} .

While the production function estimation provides estimates of Λ_{jt} and $\tilde{\mu}_{jt}$, monopsony markdowns λ_{jt} remain unobserved at this stage. As discussed in Section 2, markdowns may depend on firm size and amenities: $\lambda_{jt} = \lambda(H_{jt}, A_{jt})$. This introduces an endogeneity concern: product market rents may be correlated with markdowns, biasing inference about κ . However, firm size is observed. Therefore, using the labor wedge equation (3) in place of the wage equation (1) shifts the core identification challenge from unobserved monopsony wages Φ_{jt}^{monop} to unobserved amenities A_{jt} . This mirrors an identification challenge (IC3) faced in the rent-sharing literature.

Before addressing unobserved amenities, it is worth noting one advantage of using the labor wedge equation (3) instead of the wage equation (1): under certain conditions, unobserved amenities do not pose a challenge for identifying κ . Specifically, even when monopsony markdowns depend on firm size, if worker preferences are *multiplicatively separable* in wages and amenities, then markdowns do not depend directly on amenities. This is because, in this case, amenities shift the level of labor supply without affecting its elasticity. Therefore, unobserved amenities do not bias estimation of κ via the labor wedge equation. This multiplicative separability is a common modeling assumption for amenities (Card et al., 2018; Lamadon et al., 2022).¹⁴

However, in general, amenities may not be multiplicatively separable from wages. For this more general case, I propose a control function approach to address unobserved variation in amenities. My approach relies on the following assumption:

Assumption. The firm-specific labor supply function $H_{jt} = \mathcal{H}(\Phi_{jt}, A_{jt})$ is monotonically increasing in the value of its amenities A_{jt} , conditional on Φ_{jt} .

Under this assumption, the control function for amenities can be written as A_{jt} =

¹⁴Analogously for product markets, it is also a common modeling approach for quality or demand shifters, e.g. Hottman, Redding, and Weinstein (2016).

 $\mathcal{A}(H_{jt}, \Phi_{jt}H_{jt})$. The idea behind the amenity control function is that variation in unobserved amenities can be controlled for jointly by employment and wage bills. The intuition is that amenities shift labor supply, which breaks the direct correspondence between wages and employment in the labor supply curve. The bargaining weight κ is then identified from variation in markups, conditional on firm size in terms of employment and wage bills.

I implement equation (3) as an OLS regression of the estimated labor wedges $(\hat{\Lambda}_{jt})$ on estimated product market rents $(\hat{\mu}_{jt})$ and an approximation of the markdown function $\lambda(.)$. I approximate $\lambda(.)$ with a 4th-order polynomial in (effective) labor and wage bills. In my baseline implementation, I include sector×year fixed effects to absorb sector-specific time trends and firm fixed effects to absorb any unmodelled time invariant firm characteristics.¹⁵ Finally, I compute monopsony markdowns using the estimated bargaining parameter: $\hat{\lambda}_{jt} = \frac{1}{1-\hat{\kappa}} (\hat{\Lambda}_{jt} - \hat{\kappa} \hat{\mu}_{jt} - \hat{\epsilon}_{jt})$, where $\hat{\epsilon}$ is a residual.

My approach to estimating bargaining power is subject to several key limitations. First, measuring both labor wedges and product market rents requires estimating output elasticities, which is challenging in the context of flexible production functions, as discussed in Section 3.2. Since these elasticities enter multiplicatively into the calculations of Λ_{jt} and $\tilde{\mu}_{jt}$, any mismeasurement will mechanically bias the estimated κ upward. I refer to this as *multiplication bias*, and I address it in Appendix B.10. Second, my baseline analysis assumes that labor is a flexible input. However, if firms face labor adjustment costs, these will be absorbed into the labor wedge, potentially biasing the estimated κ . In Appendix B.11, I quantify the potential impact of labor adjustment frictions and incorporate this into a robustness check of my estimation strategy. Finally, while in one sense my markdown estimation approach is more general than existing production-based approaches—it does not impose $\kappa = 0$ —it is more restrictive in another—it imposes that all variation in markdowns reflect firm size and amenities.

4 Administrative datasets from France

Implementing the estimation approach described above requires three types of data. Firm wage premia are estimated using matched employer-employee data, which track

¹⁵While these controls are demanding, they do not eliminate identifying variation in product market rents if firm productivity and quality are not perfectly correlated. I illustrate this point in Appendix B.5.

workers over time and across different firms. Labor wedges and markups are estimated using firm balance sheet data and a survey of firm-product level output prices. I describe the data sources below and the estimation sample in Appendix A.3.

Firm balance sheet data. My source for firm balance sheet information is the *Fichier approché des résultats d'Esane* (FARE) dataset, available from 2008 to 2019. FARE is compiled by the fiscal authority of France, *Direction Générale des Finances Publiques* (DGFiP), from compulsory filings of firm annual accounting information. The dataset contains balance sheet information for all firms in France. From these datasets, I obtain information on variables such as sales, employment, material input and capital expenditure. I provide details on measurement in Appendix A.2.

Output price data. To obtain output prices at the firm level, I use the *Enquête Annuelle de Production* (EAP), available from 2009 to 2019. These are survey data compiled by the national statistical institute of France, *Institut National de la Statistique et des Études Économiques* (INSEE). The dataset contains firm-product-level sales and output for all manufacturing firms with at least 20 employees or with sales exceeding 5 million Euros, and a representative sample of manufacturing firms with less than 20 employees.

More specifically, for each year, firm, and ten-digit product code combination, the survey reports the unit of account (e.g., kilograms, meters, and pieces), total quantity sold, and total revenue. I define a product as the combination of a ten-digit product code and a unit of account, treating firms that report different units of account for the same product code as producing a different product. I drop firm-product combinations for which either quantity or revenue is missing.

COMPUTING FIRM-LEVEL OUTPUT PRICES. Let *g* denote a product, G_{jt} the set of products sold by firm *j* in year *t*, and J_{gt} the set of firms that sell product *g* in year *t*. I compute the firm-level output price as:

$$P_{jt} = \sum_{g \in G_{jt}} \mathcal{W}_{jgt} \frac{P_{jgt}}{\bar{P}_{gt}}, \text{ with } \mathcal{W}_{jgt} = \frac{P_{jgt}Y_{jgt}}{\sum_{g \in G_{jt}} P_{jgt}Y_{jgt}}, \ \bar{P}_{gt} = \sum_{j' \in J_{gt}} \left(\frac{P_{j'gt}Y_{j'gt}}{\sum_{j'' \in J_{gt}} P_{j''gt}Y_{j''gt}}\right) P_{j'gt}$$

That is, I first compute firm-product prices (P_{jgt}) for each year by dividing firm-product revenues by the corresponding quantities. I then normalize the firm-product price measure by dividing it by the sales-weighted average price of the particular product across all firms in a given year (\bar{P}_{gt}) . The firm-specific output price (P_{jt}) is then computed as the sales-weighted (W_{jgt}) average of the price index $(\frac{P_{jgt}}{P_{gt}})$ across all products sold by a given firm. This procedure follows De Ridder, Grassi, and Morzenti (2021). It is important to note that my framework does not model multiproduct firms, and therefore I do not construct theory-based firm-level price indices. Moreover, comparing output across firms is complicated by product differentiation, making price measurement inherently difficult (see De Loecker and Syverson (2021)).

Matched employer-employee data. I use annual French administrative data on employed workers from 1995 to 2018, compiled by INSEE under the *Déclarations Annuelles de Données Sociales* (DADS). These data are based on mandatory employer reports and provide worker-level information, including age, gender, earnings, hours, and occupational category. A key advantage is that hours worked are observed, enabling the construction of hourly wages and addressing concerns that earnings variation simply reflects hours variation. Observing hours also allows for more accurate measurement of firm-level labor inputs, beyond headcount employment.

I use two DADS datasets. The first is DADS-Panel, a panel of all private-sector workers born in October, which I use to estimate firm wage premia, thanks to the panel structure and employer identifiers. The second is DADS-Postes, a broader dataset covering all jobs in France but with only partial panel structure: each job appears for at most two periods under the same identifier. While this limits its use for wage premium estimation, I use DADS-Postes to k-means cluster firms, following the procedure described in the previous section, thereby maximizing the coverage of firms for which wage premia can be estimated via the DADS-Panel.

5 The Characteristics of High and Low-Wage Firms

5.1 The distribution of firm wage premia and firm market power

Firm wage premia. Table 9 reports statistics about firm wage premia. The variance is modest, accounting for 5.2% of wage dispersion, similar to the numbers for the United States, Sweden, Austria, Norway, and Italy from Bonhomme et al. (2023). Nevertheless, the dispersion of firm wage premia is a quantitatively important deviation from the law of one wage. Column 2 in Table 9 shows that a firm at the 90th percentile of

the firm wage premium distribution pays a given worker a wage that is on average 30% more than a firm at the 10th percentile. This difference amounts to approximately 4 Euros per hour or 25% of the hourly wage of the median French worker in 2016. The interquartile range is 15%, similar in magnitude to the typical estimate for the costs of job displacement (Schmieder, Von Wachter, and Heining, 2018).

Markups and labor wedges. The first row of Table 1 shows that markups are heterogeneous across firms. I estimate a median markup of 1.33. De Loecker and Warzynski (2012) estimate markups using Slovenian manufacturing data and find median markups between 1.10 and 1.28. De Loecker et al. (2020) find markups at the 75th percentile between 1.30 and 1.60 in 2016 in the US economy, while my estimates for French manufacturing is 1.60 in 2016. Edmond et al. (2023) report an interquartile range for markups of 1.31-0.97=0.34. The interquartile range for my estimates is 0.46.

The second row of Table 1 describes the distribution of labor wedges (Λ), which captures the wage-setting power of employers relative to their employees. Most French manufacturers appear to have significant wage-setting power, although France has one of the highest national minimum wages.¹⁶ Half of the firms in my sample pay less than two-thirds of workers' marginal revenue products as wages. I also find substantial dispersion of labor wedges across firms. Firms at the 75th percentile of the labor wedge distribution pay workers 76% of their marginal revenue productivity. At the 25th percentile, workers obtain half of their marginal revenue productivity.

	Mean	Median	25 th Pct	75 th Pct	Var	Var (i)	Var (ii)
Markups (μ)	1.33	1.33	1.14	1.60	0.09	0.05	0.03
Labor wedge (Λ)	0.60	0.62	0.50	0.76	0.11	0.08	0.05
Number of firms				14,342			

Table 1: Summary statistics for measured markups and labor wedges.

This table reports the summary statistics in 2016 for the measured price-cost markups and labor wedges. They are obtained by implementing my baseline approach: three-input translog production function taking into account input prices, output prices, and markup heterogeneity. Variances are reported for the log of the corresponding variable. The column Var (i) reports the variances corrected for measurement error following Krueger and Summers (1988) and Kline et al. (2020), while the column Var (ii) reports the variances between firm-groups. Markups and labor wedges are winsorized by 2%.

The set of direct estimates of labor wedges is small. I compare my estimates to

¹⁶In 2016, about 15% of French workers are earn (close to) the minimum wage. In manufacturing, the corresponding number is 10%. See Appendix A for a description of French wage-setting institutions.

those of Yeh et al. (2022) and Mertens (2020), whose estimates are methodologically the closest to mine.¹⁷ They estimate that the median US and German manufacturing firm pay 0.73 and 0.68 of the marginal revenue product of labor as wages. For the median French manufacturing firm the corresponding number is 0.62.

The role of hours and output price data. What explains the difference between my labor wedge measures for France and those for the US and Germany? As in Yeh et al. (2022) and Mertens (2020), I estimate gross output translog production functions to measure labor wedges. However, my implementation differs in several ways. First, the availability of output price data for French manufacturing firms makes it possible to: (i) measure output in terms of quantities; (ii) address unobserved input price variation; and, (iii) maintain the scalar unobservable assumption when markups and output prices are unobserved. Second, the availability of data on worker-level hours worked allows me to measure labor inputs in three ways: employment, total hours, or total effective hours (my baseline measure). Measuring labor inputs as total hours accounts for the contribution of hours to firm output, while measuring labor as total effective hours further accounts for the contribution of worker efficiency.

To assess whether methodological or measurement differences drive differences in labor wedge estimates, I compare their levels and dispersion across alternative specifications and definitions of output and labor. This comparison is shown in Table 11. I now focus on rows (1), (2), and (3), which report labor wedges obtained using a "basic" production function specification that does not address points (i), (ii), and (iii); the only difference across these rows is how labor inputs are measured.

My main finding is that the choice of labor input measure significantly affects both the level and dispersion of measured labor wedges. Using total hours yields lower labor wedges with substantially less dispersion than using employment. When labor is measured as total effective hours, the median labor wedge is 0.64; it rises to 0.68 with total hours and to 0.71 with employment. This difference in labor input measurement explains much of the gap between my labor wedge estimates and those of Yeh et al. (2022) and Mertens (2020). Furthermore, the 90-10 labor wedge difference is much

¹⁷However, the interpretation of these labor wedges differ between my paper and those in Yeh et al. (2022) and Mertens (2020). In the aforementioned papers, labor wedges are interpreted as monopsony markdowns, corresponding to the case where workers do not have bargaining power ($\kappa = 0$) in my framework in Section 2. Allowing $\kappa > 0$ also helps rationalize that about 7% of firms have a labor wedge greater than 1. Indeed, firms with $\Lambda > 1$ tend to have much higher markups.

larger when using employment (0.70) instead of effective hours (0.60).

Similarly, both the level and dispersion of measured markups fall when I measure labor inputs as total hours instead of employment. When measured as employment, the median markup is 1.41 and 90-10 difference is 1.14. Instead, when measured as total hours, the median markup is 1.39 and 90-10 difference is 1.10. When measured as total effective hours, the corresponding numbers fall further to 1.37 and 1.05.

These findings suggest that using employment instead of hours or effective hours understates the *level* of firm labor market power and overstates their *dispersion*. This distinction is important because recent models of labor and product market power emphasize that the level of labor wedges or markups acts as a uniform tax on labor demand across firms, while dispersion leads to misallocation of labor (Berger et al., 2022; Edmond et al., 2023). Since data on hours are not consistently available across countries, these findings imply that researchers should interpret the estimated moments of labor wedges and markups with caution. In Appendix B.4, I discuss two reasons for why using employment in place of total hours leads to upward-biased measures of labor wedges: (i) omitted variable bias (average hours); and (ii) mismeasurement of output elasticities, even with known production function parameters.

The availability of output price data also enables the estimation of richer specifications that account for output and input price biases, as well as potential violations of scalar unobservability. I discuss these additional results in Appendix B.4.

5.2 Prices, markups, and labor wedges among high-wage firms

I now present empirical patterns between firm wage premia and firm market power that are not readily explained by existing labor market monopsony models where firms operate in competitive product markets and do not bargain.

To document how output prices, markups, and labor wedges vary across highwage and low-wage firms, I separately regress log output prices (p_{jt}), markups (μ_{jt}), and labor wedges (Λ_{jt}) on deciles of firm wage premia (ϕ_{jt}), controlling for 5-digit sector×year fixed effects, and TFPQ. The estimated coefficients are shown in Figure 1. I find similar patterns when I do not control for differences in TFPQ (see Figure 24). Figure 25 also shows similar patterns when I compare by firm size (sales).

Figure 1 shows that high-wage firms tend to charge higher output prices and markups.

While theory suggests that more productive firms charge higher markups (Atkeson and Burstein, 2008) and pay higher wages (Card et al., 2018), productivity alone cannot explain these patterns: more productive firms should charge *lower* output prices, contrary to the data. Indeed, the positive correlation between prices, markups, and wage premia persists even after conditioning on productivity, pointing to differences in product quality as a key driver of wage premia.



Figure 1: Prices, markups, and labor wedges by high-wage status, conditional on TFP. Notes: This figure shows how log output prices, price-cost markups, and labor wedges vary by deciles of firm wage premia relative to firms in the first decile, controlling for TFP(Q) and 5-digit sector×year fixed effects. Decile 10 represents high-wage firms. Confidence intervals at the 95% level are plotted.

The figure also shows that high-wage firms pay a larger share of the marginal revenue product of labor as wages (i.e., higher labor wedges). This pattern is inconsistent with monopsonistic models where workers have no bargaining power, in which the labor wedge coincides with the monopsony markdown and typically declines with wages (Burdett and Mortensen, 1998; Berger et al., 2022; Lamadon et al., 2022).¹⁸ However, as equation (3) shows, when workers have bargaining power ($\kappa > 0$), labor wedges increase with markups—consistent with the observed empirical patterns.

5.3 Worker bargaining power & firm monopsony markdowns

Worker bargaining power. My estimates, reported in Table 2, suggest that worker bargaining power is relatively low: workers obtain around 12% of firm profits.¹⁹

¹⁸The model in Berger et al. (2022) allows monopsony markdowns to rise with wages when jobs across labor markets are closer substitutes than jobs within the same labor market.

¹⁹Table 25 in Appendix D reports these estimates by sector. The sector-specific bargaining power are estimated using my baseline regression specification (following column (4) in Table 2). In all 2-digit French manufacturing sectors, I find that worker bargaining power is below 0.3.

Column (1) presents the pooled OLS estimate of bargaining power, yielding an estimate of approximately 0.14. Under the following conditions (also mentioned in Section 3.3), this is an unbiased estimate of κ : (i) monopsony markdowns are homogenous across firms ($\lambda_{jt} = \lambda_t$); (ii) the labor wedges do not reflect labor adjustment costs; and (iii) no mismeasurement of labor wedges and product market rents. Column (2) additionally includes firm fixed effects, and returns a similar estimate of around 0.13. Although this specification continues to rely on conditions (ii) and (iii), it relaxes condition (i), allowing for firm-specific monopsony markdowns ($\lambda_{jt} = \lambda_j$) as long as these markdowns do not exhibit firm-specific trends over time.

	Labor wedge (Λ_{it})						
	(1)	(2)	(3)	(4)	(5)	(6)	
Product market rents ($\tilde{\mu}_{jt}$)	0.136	0.126	0.136	0.125	0.131	0.071	
	(0.007)	(0.009)	(0.007)	(0.009)	(0.006)	(0.036)	
Sector×year fixed effects	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	
Firm fixed effects		\checkmark		\checkmark		\checkmark	
Control for $\lambda(.)$			\checkmark	\checkmark	\checkmark	\checkmark	
IV for measurement error					\checkmark	\checkmark	
Observations	102.777	99.791	102.777	99.791	79,424	76.747	

Table 2: Estimated worker bargaining power.

This table reports the estimated workers' bargaining power parameter. Even columns controls for firm fixed effects. Columns (3) through (6) controls for differences in monopsony markdowns reflecting differences in amenities. Columns (5) and (6) instruments measured product market rent $\tilde{\mu}_{jt}$ with its lag to address classical measurement error. Bootstrapped standard errors are reported in parentheses.

Columns (3) and (4) implement the control function approach described in Section 3.3 to address unobserved amenities. These specifications accommodate the case where amenities are proportional labor supply shifters and do not directly determine markdowns (see the discussion on multiplicative separability in Section 3.3), but also the more general case where amenities may directly affect markdowns. I continue to find a bargaining power estimate of around 0.12 to 0.14, suggesting that amenities do not have a significant direct impact on monopsony markdowns.

Finally, columns (5) and (6) further instrument measured product market rents using their lags. As outlined in Section 3.3 and discussed in Appendix B.10, measurement error in labor wedges and product market rents will be positively correlated since they depend on the same variables. This correlation could mechanically bias the bargaining power estimates upward. Under a classical measurement error assumption, using lagged product market rents as instruments corrects this bias. These specifications return estimates between 0.13 and 0.07, which align with those from the rent-sharing literature, typically finding values between 0.05 and 0.15 (Card et al., 2018; Jäger et al., 2020). However, studies using external, innovation-based instruments tend to find larger values, around 0.30 (Van Reenen, 1996; Kline et al., 2019).

It is important to note two caveats that I address in detail in Appendix B. First, the inclusion of firm fixed effects does not adequately address the presence of labor adjustment costs, which are unobserved, time-varying, likely correlated with product market rents, and directly determine labor wedges. I address this concern in Appendix B.11. Second, when output elasticities are not well-identified, any biases that arise in their estimation will be serially correlated, mechanically causing an upward bias in $\hat{\kappa}$. In this case, using lagged product market rents as instruments does not adequately address the mechanical upward bias. I address this concern in Appendix B.10.

Monopsony markdowns. I compute the implied monopsony markdowns taking the bargaining power estimate from specification (4) in Table 2 as my baseline estimate.²⁰ Table 3 reports the estimated monopsony markdowns. At the median, the markdown is 45% of the MRPL, implying significant monopsony power. The corresponding firm-specific labor supply elasticities are 0.54, 0.82, and 1.33 at the 25th, 50th, and 75th percentiles. These estimates are similar to those found for the US labor market using the Burdett-Mortensen model by Webber (2015), who reports firm-specific labor supply elasticities of 0.44, 0.75, and 1.13 at the same percentiles. In an oligopsonistic setting, Berger et al. (2022) find labor supply elasticities ranging from 0.76 to 3.74 in the US. Meanwhile, a meta-analysis by Sokolova and Sorensen (2020), covering over 700 studies, finds a median labor supply elasticity of around 1.10. My median markdown-implied labor supply elasticity (0.82) is somewhat lower than this benchmark.

There is a significant difference between the median labor wedge (0.62) and median monopsony markdown (0.45), which is accounted for by bargaining power and markups. The production approach to estimating markdowns yields estimates of monopsony markdowns only when workers have no bargaining power. When bar-

²⁰Since the monopsony markdowns are measured after estimating worker bargaining power, their accuracy depends on the extent to which it is possible to obtain unbiased and precise estimates of κ . I compare the inferred markdowns under different values of κ in Appendix B.10.

gaining power is positive, these estimates instead reflect labor wedges. Consequently, production-based markdown estimates that do not account for bargaining power understate the true extent of firm monopsony power.

Summary statistics	Mean	Median	25 th Pct	75 th Pct	Var	Var (i)	Var (ii)
Markdowns (λ_{jt})	0.46	0.45	0.35	0.57	0.04	0.02	0.01

Table 3: The distribution of monopsony markdowns.

This table reports the summary statistics in 2016 for the estimated monopsony markdowns. Markdowns are computed using the estimated bargaining parameter κ from specification (4) in Table 2. The column Var (i) reports the variances corrected for measurement error following Krueger and Summers (1988) and Kline et al. (2020), while the column Var (ii) reports the variances for firmgroups. Each variable is winsorized by 2%.



Figure 2: Monopsony markdowns by high-wage status and firm size.

How do monopsony markdowns differ across firms and over time? Figure 2 shows how markdowns vary by deciles of firm wage premia and employment. I find that markdowns increase with firm wage premia, although this gradient is quite flat. When I compare markdowns by firm size, I find that larger firms have lower markdowns larger firms have more monopsony power. The negative relationship between markdowns and firm size is consistent with the predictions of monopsonistic or oligopsonistic labor market models in which markdowns vary endogenously, such as Gouin-Bonenfant (2022), Berger et al. (2022), and Jarosch, Nimczik, and Sorkin (2024). Overall, Figure 27 shows that there are no significant trends observed, though markdowns show a slight increase over the sample period, suggesting a slight increase in labor market competition.

Notes: The figures shows how monopsony markdowns vary by deciles of firm wage premia and employment relative to firms in the first decile, controlling for 2-digit sector×year fixed effects. Decile 10 represents high-wage or large firms. Confidence intervals are at the 95% level.

6 A General Equilibrium Model of Firm & Worker Power

The previous section provides descriptive evidence on firm market power and worker bargaining power. However, it leaves several key questions unanswered: What are the sources of firm heterogeneity that shape the observed distributions of wages and market power? How important are markups in determining wages? And, can strengthening worker bargaining power lead to welfare improvements?

Answering these questions requires specifying the sources of product and labor market power in the model in Section 2. I assume that product varieties and workplace amenities are both horizontally and vertically differentiated, leading to firm market power in these markets. The specific functional form that I adopt for consumption and labor supply preferences are both nested-CES. I assume that product and labor market structures are oligopolistic and oligopsonistic, giving rise to firm-specific markups and markdowns that depend on their market shares (Atkeson and Burstein, 2008; Berger et al., 2022). Firms differ in productivity, quality, and amenities.²¹

6.1 Environment

There is a continuum of product markets $g \in [0, G]$ and local labor markets $s \in [0, S]$. Each product and local labor market contains an exogenously given finite number of firms n_g and n_s . Each firm j belongs to one product market and one local labor market.

Labor supply. The representative household maximizes utility by choosing the amount of numeraire final goods to consume *C* and effective labor to supply to each firm H_{is} :

$$U = \max_{\{C,H_{js}\}} C - \frac{H^{1+\varphi}}{1+\varphi}$$

subject to the budget constraint $C = \Phi H + \Pi$, where Φ is the aggregate wage index and Π are aggregate profits. Jobs across and within markets are imperfect substitutes. Aggregate labor supply H is a composite of market-level labor supply H_s , with a constant elasticity of substitution ν , governing the degree of horizontal job differentiation across markets. Within each market, the market-level labor supply is a composite of firm-specific labor supply with a constant elasticity of substitution η , governing the

²¹Appendix C.7 shows that CES-implied measures of quality and amenities are closely aligned with those implied by variable elasticity of substitution (VES) preferences, which nests CES as a special case.

degree of horizontal job differentiation within markets. The aggregate and marketlevel labor supply aggregators are:

$$H = \left[\int_0^S \tilde{A}_s^{-\frac{1}{\nu}} H_s^{\frac{\nu+1}{\nu}} \, ds \right]^{\frac{\nu}{\nu+1}} \quad \text{and} \quad H_s = \left[\sum_{j=1}^{n_s} \tilde{A}_{js}^{-\frac{1}{\eta}} H_{js}^{\frac{\eta+1}{\eta}} \right]^{\frac{\eta}{\eta+1}}$$

where \tilde{A}_s are market-level labor supply shifters such that $\int_s \tilde{A}_s ds = 1$, and \tilde{A}_{js} are firm-specific labor supply shifters such that $\sum_{j=1}^{n_s} \tilde{A}_{js} = 1$. When $\eta \ge \nu$, jobs within markets are closer substitutes than jobs across markets. When $\eta \to \infty$ ($\nu \to \infty$), there is no job differentiation within (between) markets. The firm-specific labor supply curve:

$$H_{js} = \tilde{A}_{js} \tilde{A}_s \Phi^{\eta}_{js} \Phi^{\nu-\eta}_s \Phi^{-\nu} H \tag{8}$$

I refer to $A_{js} \equiv \tilde{A}_{js}\tilde{A}_s$ as amenities. This provides a microfoundation for the labor supply curve in Section 2, $H_{js} = \mathcal{H}_s(\Phi_{js}, A_{js})$, where the subscript *s* on the function $\mathcal{H}_s(.)$ accommodates the endogenous market-wide and economy-wide aggregates.

Product demand. The competitive final good producer produces good *Y* by combining the output of firms in each product market. The CES goods aggregator across markets and within markets are:

$$Y = \left[\int_0^G \tilde{D}_g^{\frac{1}{\theta}} Y_g^{\frac{\theta-1}{\theta}} dg \right]^{\frac{\theta}{\theta-1}} \quad \text{and} \quad Y_s = \left[\sum_{j=1}^{n_g} \tilde{D}_{jg}^{\frac{1}{\sigma}} Y_{jg}^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}$$

where \tilde{D}_g are market-level demand shifters such that $\int_g \tilde{D}_g dg = 1$, and \tilde{D}_{jg} are firmspecific demand shifters such that $\sum_{j=1}^{n_g} \tilde{D}_{jg} = 1$. The within-market and betweenmarket elasticities of substitution between goods are σ and θ . When $\sigma \geq \theta$, goods within a market are more substitutable than goods across markets. When $\sigma \to \infty$ $(\theta \to \infty)$, there is no product differentiation within (between) markets. Firms face the following inverse goods demand curve:

$$Y_{jg} = \tilde{D}_{jg}\tilde{D}_g P_{jg}^{-\sigma} P_g^{\sigma-\theta} Y$$
⁽⁹⁾

where I refer to $D_{jg} \equiv \tilde{D}_{jg}\tilde{D}_g$ as product quality. This provides a microfoundation for the product demand curve in Section 2, $Y_{jg} = \mathcal{G}_g(P_{jg}, D_{jg})$, where the subscript *g* on the function $\mathcal{G}_g(.)$ accommodates the endogenous market- or economy-wide aggregates.

Resource constraint. Final goods can be consumed by the representative household or used as factor inputs—capital and material—such that Y = C + K + M, where

 $K = \int_g \sum_j^{n_g} K_{jg} \, dg$ and $M = \int_g \sum_j^{n_g} e^{\tau_{jg}} M_{jg} \, dg$ are aggregate capital and material use.²² Firm level production technologies are $Y_j = \Omega_j K_j^{\alpha_k} M_j^{\alpha_m} H_j^{\alpha_h}$.

Wage bargaining. The wage bargaining problem is as described in Section 2.

6.2 The Social Planner's Equilibrium

In this section, I characterize the distortions caused by firm market power and show how worker bargaining power can improve social efficiency. The social planner chooses capital, materials, and labor at each firm to maximize household utility, subject to the same preferences, production technologies, and resource constraint as the market economy. I write the planner's problem formally in Appendix C.4.

Definitions. Let μ denote the common price-cost markup and λ the common monopsony markdown that would prevail if firms had identical market power. Define the common labor wedge (Λ) as the labor wedge that would arise under this common markup and markdown: $\Lambda \equiv \kappa \tilde{\mu} + (1 - \kappa)\lambda$, where $\tilde{\mu} \equiv \left(1 - \frac{\alpha_k + \alpha_m}{\mu}\right) \frac{\mu}{\alpha_h}$ represents the corresponding product market rents.

Social efficiency under firm market power. To understand how firm product and labor market power distort social efficiency, I compare the *aggregate* labor demand conditions in the planner's economy and the market economy:

Planner's equilibrium	Market equilibrium			
$H^{\varphi} = \alpha_h \frac{Y}{H}$	$H^{arphi} = rac{\Lambda}{\mu} \mathbf{\Theta} lpha_h rac{Y}{H}$			

This comparison shows that the planner's choice of labor allocation equalizes the marginal disutility of labor supply and the marginal product of labor. In contrast, the market allocation deviates from this benchmark due to two wedges: a *uniform tax* $(\frac{\Lambda}{\mu})$ and a *misallocation tax* (Θ), both of which depend on κ .

²²Material price heterogeneity helps the model to match the empirical firm size distribution. There is broad evidence that material prices vary across firms due to differences in quality. For example, the literature on input trade liberalization shows that firms use inputs of different quality and that tariff reductions lead some firms to import higher quality inputs, raising their productivity and size (Kugler and Verhoogen, 2012; Fieler, Eslava, and Xu, 2018). Material quality heterogeneity can be accommodated by my model; the estimation approach in Section 3 accounts for material input quality heterogeneity by using output price data, following De Loecker et al. (2016). For the same reason, my model can also accommodate capital quality heterogeneity, but I do not model this explicitly. However, not all sources of firm heterogeneity can be accommodated by my model. For example, my model cannot accommodate heterogeneous monopsony power in material and capital markets due to the absence of input price data. More detailed discussions are available in Section 3.

The uniform tax reflects how firm market power proportionally lowers labor demand across all firms. This result generalizes existing findings: Edmond et al. (2023) show this for markups (μ), and Berger et al. (2022) show it for monopsony markdowns (λ). I extend these insights to environments with worker bargaining power: the common labor wedge (Λ)—hence, the uniform tax—depends on κ .

The misallocation tax shows that dispersion in markups and markdowns—thus, labor wedges—lead to *misallocation of labor* across firms. It is defined as:

$$\boldsymbol{\Theta} \equiv \int_{g} \left[\sum_{j \in g} \left(\frac{\Lambda_{jg} / \Lambda}{\mu_{jg} / \mu} \right) \frac{P_{jg} Y_{jg}}{P_{g} Y_{g}} \right] \frac{P_{g} Y_{g}}{PY} \, dg$$

The distortion term $\frac{\Lambda_{ig}/\Lambda}{\mu_{ig}/\mu}$ captures firm-level deviations from the common labor wedge and markup. These deviations are weighted by firm size, measured by revenue shares. When larger firms face lower labor wedges and/or charge higher markups, then $\Theta <$ 1, indicating misallocation and lower aggregate labor demand. If all firms share the same markup and markdown, then $\Theta = 1$, so labor is efficiently allocated across firms.

Conditions for social efficiency. Under what conditions do the market's and the planner's aggregate labor demand conditions coincide? The comparison above shows that they coincide when (i) firms are price-takers in both product and labor markets ($\mu_{jg} = 1$ and $\lambda_{jg} = 1$), and (ii) workers are wage-takers ($\kappa = 0$, so $\Lambda_{jg} = 1$). Under these conditions, the uniform tax and the misallocation tax are both equal to one.

Worker bargaining power offsets the uniform tax. I now show that worker bargaining power can help close the gap between the market and planner equilibria by counteracting the effect of firm market power as a uniform tax on aggregate labor demand. Following from the definitions above, the uniform tax can be expressed as:

$$\frac{\Lambda}{\mu} = \kappa \, \frac{\widetilde{\mu}}{\mu} + (1 - \kappa) \, \frac{\lambda}{\mu}$$

This expression shows that the uniform tax is a weighted average of two polar cases. When $\kappa = 0$, it reflects the full effect of firms' common component of market power, $\frac{\lambda}{\mu}$, a "double market power" effect, as termed by Kroft et al. (2024). When $\kappa = 1$, workers behave as monopolist wage-setters and fully capture the rents generated by firms' product market power.

For values of $\kappa \in [0,1]$, there exists a level of worker bargaining power that ex-

actly offsets the uniform tax: $\frac{\Lambda}{\mu} = 1$ when $\kappa = \bar{\kappa} \equiv \frac{\mu - \lambda}{\bar{\mu} - \lambda} \in [0, 1]$. At this threshold, the opposing effects of firms as monopsonist wage-setters and workers as monopolist wage-setters exactly cancel out. Moreover, $\bar{\kappa}$ decreases with the common markup μ , implying that it is easier to eliminate the uniform tax when markups are high—since higher markups increase the rents available for workers to extract.

Worker bargaining power offsets the misallocation tax. I now show that worker bargaining power can also reduce the misallocation of labor caused by dispersion in monopsony markdowns (i.e., $\lambda_{js} \neq \lambda$). To isolate this effect, I assume for now that firms charge a common markup, so $\mu_{jg} = \mu$. The misallocation tax then simplifies to:

$$\Theta|_{\mu_{jg}=\mu} = \int_{g} \left[\sum_{j \in g} \left(\frac{\Lambda_{jg}}{\Lambda} \right) \frac{P_{jg} Y_{jg}}{P_{g} Y_{g}} \right] \frac{P_{g} Y_{g}}{PY} \, dg \quad \text{where} \quad \frac{\Lambda_{jg}}{\Lambda} = \frac{\kappa \widetilde{\mu} + (1-\kappa)\lambda_{js}}{\kappa \widetilde{\mu} + (1-\kappa)\lambda}$$

When larger firms exert more monopsony power (i.e., $cov(\lambda_{js}, P_{jg}Y_{jg}) < 0$), the misallocation tax becomes distortionary: $\Theta|_{\mu_{jg}=\mu} < 1$. This reflects that large firms employ too little labor relative to the planner's solution. However, with full worker bargaining power ($\kappa = 1$), the labor wedge becomes uniform across firms despite heterogeneity in markdowns: $\frac{\Lambda_{jg}}{\Lambda}|_{\kappa=1} = 1$. In this case, bargaining neutralizes the allocative inefficiencies induced by dispersion in monopsony power, i.e., $\Theta|_{\mu_{ig}=\mu,\kappa=1} = 1$.

When markups are also dispersed ($\mu_{jg} \neq \mu$), the effect of worker bargaining power on allocative efficiency becomes ambiguous. To illustrate this, consider the case where larger firms not only exert more monopsony power but also charge higher markups: $cov(\lambda_{js}, P_{jg}Y_{jg}) < 0$ and $cov(\mu_{jg}, P_{jg}Y_{jg}) > 0$.

When workers have no bargaining power, the distortion term in the misallocation tax becomes $\frac{\Lambda_{jg}/\Lambda}{\mu_{jg}/\mu}\Big|_{\kappa=0} = \frac{\lambda_{js}/\lambda}{\mu_{jg}/\mu}$, which is negatively correlated with firm size. In this case, $\Theta|_{\kappa=0} < 1$, and large firms are too small relative to the planner's choice.

In contrast, when workers have full bargaining power, the distortion term becomes $\frac{\Lambda_{jg}/\Lambda}{\mu_{jg}/\mu}\Big|_{\kappa=1} = \frac{1-(\alpha_k+\alpha_m)\mu_{jg}^{-1}}{1-(\alpha_k+\alpha_m)\mu^{-1}}$, which is positively correlated with firm size. In this case, $\Theta|_{\kappa=0} > 1$, implying that large firms too large relative to the planner's choice. This is because, when workers fully extract the rents from product market power, firms are compelled to over-produce.

Limits of worker bargaining power in achieving full social efficiency. Although worker bargaining power can mitigate distortions to labor demand, it cannot achieve full efficiency. This is because bargaining only directly affects labor demand—through

labor wedges—but markups also directly distort the demand for capital and materials. Therefore, the presence of markups generates inefficiencies that bargaining power alone cannot correct.

6.3 Calibrating the model

I now describe my calibration approach. The calibrated parameter values are reported in Table 4. I first calibrate parameters that govern markups and markdowns (θ , σ , ν , η), then measure product quality (D_{ig}) and amenities (A_{is}).

Parameters related to markups. The relationship between markups and product market shares in this model can be written as:

$$\rho_{jg}^{-1} \equiv \frac{\mu_{jg} - 1}{\mu_{jg}} = \frac{1}{\sigma} + \left(\frac{1}{\theta} - \frac{1}{\sigma}\right) \frac{P_{jg} Y_{jg}}{\sum_{j'}^{n_g} P_{j'g} Y_{j'g}}$$

I calibrate the within-market elasticity of substitution (σ) to match the level (median) of estimated markups, and the between-market elasticity of substitution (θ) to match the passthrough of product market shares to estimated markups. Product market shares are measured as sales shares within 5-digit sectors.

Parameters related to markdowns. The relationship between oligopsonistic markdowns and wage bill shares in this model can be written as:

$$\xi_{js} \equiv \frac{\lambda_{js}}{1 - \lambda_{js}} = \eta + (\nu - \eta) \frac{\Phi_{js} H_{js}}{\sum_{j'}^{n_s} \Phi_{j's} H_{j's}}$$

I calibrate the within-market elasticity of substitution (η) to match the level (median) of estimated markdowns, and the between-market elasticity of substitution (ν) to match the passthrough of labor market shares to estimated markdowns. Labor market shares are measured as wage-bill shares within 5-digit sectors×commuting-zone pairs.

Appendix C.6.1 provides further detail on the calibration of the markup/downrelated parameters and discusses model fit. Appendix C.6.2 discusses the measurement issues surrounding market shares.

Calibrating heterogeneous productivity, quality, and amenities. Firm productivity (TFPQ) is estimated directly from the data in Section 3. Given the parameters governing product demand and labor supply curves, I back out firm heterogeneity in product quality and amenities. The product and labor market shares of a given firm are:

$$\frac{P_{jg}Y_{jg}}{P_gY_g} = D_{jg} \left(\frac{P_{jg}}{P_g}\right)^{1-\sigma} \quad \text{and} \quad \frac{\Phi_{js}H_{js}}{\Phi_sH_s} = A_{js} \left(\frac{\Phi_{js}}{\Phi_s}\right)^{1+\eta}$$

Using these equations to measure quality and amenities requires measuring firmlevel output prices and (hourly) wage premia. A key advantage of the French administrative data over many existing datasets (e.g., the US LEHD and Census data) is that it allows one to do so, instead of relying on earnings data, for example.

I compute D_{jg} such that $\sum_{j}^{n_g} D_{jg} = 1$ and A_{js} such that $\sum_{j}^{n_s} A_{js} = 1$. Similarly, the sectoral product demand and labor supply shifters (D_g and A_s) are measured using the relative size of sectors: $\frac{P_g Y_g}{PY} = D_g (\frac{P_g}{P})^{1-\theta}$ and $\frac{\Phi_s H_s}{\Phi H} = A_s (\frac{\Phi_s}{\Phi})^{1+\nu}$. Appendix C.7 compares the quality and amenity measures under non-CES preferences that nest the CES as a special case, finding that CES-implied measures are closely aligned with non-CES-implied measures.

Parameter		Value	Source
Frisch labor supply elasticity	φ	0.25	Chetty (2012)
Between-market e.o.s. (labor)	ν	0.67	Calibrated
Within-market e.o.s. (labor)	η	0.97	Calibrated
Between-market e.o.s. (product)	θ	1.21	Calibrated
Within-market e.o.s. (product)	σ	5.16	Calibrated
Worker bargaining power	κ	0.12	Estimated (see Section 3.3)
Number of firms within goods markets	n_g	[71, 28, 125]	FARE-EAP sample
Number of firms within labor markets	n_s	[2, 1, 3]	FARE-EAP sample
Labor elasticity of output	α_h	0.47	Average of estimates
Material elasticity of output	α_m	0.47	Average of estimates
Capital elasticity of output	α_k	0.06	$1 - \alpha_h - \alpha_m$
Log TFPQ	ω_{j}	[0.00, -0.07, 0.89]	Production function estimation
Log product quality	d_i	[0.00, 0.08, 3.80]	Sales shares
Log non-wage amenities	a_{i}	[0.00, -0.04, 0.83]	Wage bill shares
Log material price heterogeneity	τ_i	[0.00, -0.16, 2.96]	Firm size distribution (eff. labor)

Table 4: Calibrated parameters.

Notes: This table reports the calibrated values of the model parameters. The abbreviation "e.o.s." means "elasticity of substitution". The numbers in the brackets represent the mean, median, and standard deviation, respectively. TFPQ, quality, amenities, and material price heterogeneity have been de-meaned.

Other parameters. The parameter value for the Frisch labor supply elasticity is obtained from Chetty (2012). I calibrate workers' bargaining power to 0.12 using my estimates in Table 2 from specification (4). The number of firms within each 5-digit sector is obtained from my FARE-EAP sample. The production function parameters are calibrated as the average of my production function estimates. Finally, material price heterogeneity τ_i is calibrated to match the firm size distribution (in efficiency units of
labor). Combined with the calibrated values of non-wage amenities, the model exactly reproduces the empirical firm wage premium distribution.

7 Quantitative Findings

In this section I use the calibrated model for three main exercises. First, I describe how the measured product quality and amenities vary across high-wage and low-wage firms. Second, I decompose the passthrough of firm heterogeneity to wage premia into various channels and quantify their contributions. Third, I quantify the importance of firm market power and worker bargaining power for wages and welfare.

7.1 Product quality and amenities among high-wage firms

Analogously to Section 5, I compare (log) product quality and amenities across deciles of firm wage premia, controlling for 5-digit sector × year fixed effects. Figure 26 in Appendix D shows that the difference between high-wage and low-wage firms in quality and amenities are larger than the difference in productivity (TFPQ). While firms in the top decile of the wage premium distribution are on average 15% more productive than firms at the bottom decile, they have over 100% greater quality. That is, for a given price, firms in the top decile can sell over 100% more goods than firms in the bottom decile. Amenities also increase slightly more steeply with firm wage premia than productivity, consistent with the findings of Lamadon et al. (2022) for US firms.

7.2 On the passthrough of firm heterogeneity to wages

In this section, I show that the degree of product differentiation governs the extent of productivity passthrough relative to product quality passthrough. I also show how firm size dampens the passthrough of firm-specific shocks, leading to a smaller passthrough among large, high-wage firms. Finally, I show that product quality is as important as productivity in explaining firm wage premia.

How firm heterogeneity affects firm wage premia. I write wage premia (in *logs*) from the model of Section 6 as:

$$\phi_{j} = \underbrace{\frac{(\sigma-1)}{1+\eta+\alpha_{h}(\sigma-1)}\omega_{j} + \frac{1}{1+\eta+\alpha_{h}(\sigma-1)}(d_{jg}-a_{js}) - \frac{\alpha_{m}(\sigma-1)}{1+\eta+\alpha_{h}(\sigma-1)}\tau_{j}}_{\text{direct effect}} + \underbrace{\frac{1+\alpha_{h}(\sigma-1)}{1+\eta+\alpha_{h}(\sigma-1)}\log\left(\kappa\widetilde{\mu}_{jg}+(1-\kappa)\lambda_{js}\right)}_{\text{variable market power effect}} - \frac{\sigma}{1+\eta+\alpha_{h}(\sigma-1)}\log\mu_{jg}}_{\text{variable market power effect}}$$
(10)
$$+\underbrace{\frac{(\sigma-\theta)p_{g}+(\eta-\nu)\phi_{s}}{1+\eta+\alpha_{h}(\sigma-1)}}_{\text{firm size effect}} + \underbrace{\frac{1}{1+\eta+\alpha_{h}(\sigma-1)}\log\left(\frac{\gamma}{\Phi^{-\nu}H}\right)}_{\text{general equilibrium effect}} + \underbrace{\frac{\alpha}{\alpha}}_{\text{constant}}$$

The first line of equation (10) shows the *direct effect* of firm heterogeneity on wage premia, reflecting its impact on a firm's labor demand and/or supply. A higher labor output elasticity α_h reduces the passthrough of productivity, product quality, and amenities, since a smaller increase in employment can now achieve a given increase in output. When jobs within markets are closer substitutes (higher η), the passthough of firm heterogeneity becomes smaller, because a smaller wage increase is sufficient to achieve a given increase in employment.

The substitutability of product varieties (σ) affects the passthrough of productivity and product quality in opposite directions. When σ is higher, productivity passthrough increases, while quality passthrough decreases. Intuitively, with higher σ , product varieties become closer substitutes, making consumers more price-sensitive. This magnifies the impact of productivity heterogeneity on labor demand and wages, making productivity a more significant determinant of wage premia in high- σ environments.

The second line shows that endogenous changes in labor wedges and markups in response to firm-specific shocks also affect wage premia. I refer to this as the *variable market power effect*. A higher markup can raise wage premia through *rent-sharing*—by increasing the labor wedge—with the magnitude of this effect depending on the level of worker bargaining power, κ . However, a higher markup can also reduce wage premia by inducing the firm to *restrict output* and reduce labor demand. The passthrough elasticity of firm-specific shocks to wages thus depends on the relative strength of these opposing forces.

In addition, labor wedges depend on monopsony markdowns. A positive firmspecific productivity shock raises firm size, increasing the firm's monopsony power and widenening its markdown. This leads to incomplete passthrough of the shock to wage premia, as shown by Berger et al. (2022). In my model, the magnitude of this markdown effect is attenuated by worker bargaining power: higher κ reduces the sensitivity of labor wedges to monopsony power, thereby increasing wage passthrough.

The third line shows that firm wage premia depend on market-level price and wage indices. This *firm size effect* arises when firms are large relative to the markets they belong (i.e., they are not atomistic). When firms are large, their price and wage-setting actions have spillover effects on their competitor's actions. For example, if a large firm reduces prices, it lowers the market-level price index, forcing its competitors to also reduce prices. Since $\sigma > \theta$, a lower sectoral price index induces firms to produce and hire less, reducing wages. Similarly, given that $\eta > \theta$, a higher sectoral wage index forces firms to pay higher wages.

The final term reflects the general equilibrium effects on wage levels, which impact all firms equally and thus do not affect the distribution of firm wage premia. The numerator, aggregate output *Y*, captures the effect of aggregate product demand on wages, while the denominator, $\Phi^{-\nu}H$, represents an endogenous aggregate labor supply shifter (see equation (8)). Together, these general equilibrium effects capture the role of markups and markdowns in distorting aggregate labor demand.

Quantitative results. I now present the passthrough elasticities of firm wage premia with respect to firm-specific productivity, quality, and amenity shocks. To compute the passthrough of a shock, I randomly select one firm from each sector and assign it a 1% positive shock, while holding general equilibrium aggregates constant. I then compute the average elasticity of wages to the shock across shocked firms. Panel (A) of Table 5 presents the passthrough elasticities.

The passthrough of a productivity shock is 0.99, around four times larger than the passthrough of a product quality shock, as the last column of Panel (A) shows. Columns one through five decomposes this difference, showing that the *direct effect* accounts for most it. Given the calibrated value of σ , the direct effect is larger for productivity passthrough. However, for sufficiently low σ , the passthrough of product quality can be higher than that of productivity, as Table 26 in Appendix C shows.

The *variable market power* channel also plays a role in the passthrough of productivity shocks. The first row of columns two and three in Table 5 shows how variable markups and markdowns contribute to the productivity passthrough. A higher productivity leads to higher markups, leading to an increase in rent-sharing (higher labor wedges) but also a stronger incentive to restrict output. Overall, the output restriction effect dominates, implying that increased markups dampen the passthrough.

Panel (A)			Total			
	Direct	Var. mar	Firm	size	passthrough	
		$\Delta \log \Lambda$	$-\Delta \log \mu$	Δp_s	$\Delta\phi_s$	
Δ Productivity	1.07	0.02	-0.08	-0.07	0.06	0.99
Δ Product quality	0.26	0.01	-0.02	-0.02	0.01	0.24
Δ Amenity	-0.26	0.01	-0.02	-0.02	0.01	-0.28
Panel (B)	Firm	n wage pre	emium dist	ributior	۱	
	10 th pct.	50^{th} pct. 90^{th}		pct.		
Δ Productivity	1.06	1	07	0.78		
Δ Product quality	0.25	().26	0.	19	
Δ Amenity	-0.26	-().26	-0.	32	

Table 5: Passthrough of firm heterogeneity to firm wage premia.

This table reports the passthrough of firm heterogeneity to firm wage premia. The passthrough measures are with respect to a positive 1% shock. Panel (A) decomposes the total passthrough into different effects. The final column presents the total passthrough, which is the sum of columns 1 through 5. Panel (B) looks at the passthrough at different percentiles of the firm wage premium distribution.

The *firm size* channel may further amplify or dampen the passthrough of productivity shocks. Given the calibrated σ , the reduction in sectoral output prices reduces the passthrough elasticity by -0.07. However, given the calibrated substitutability of jobs within sectors η , the response of the market-level wage index is 0.06, amplifying the productivity passthrough.

The total passthrough of an amenity shock is larger than that of a product quality shock. Table 5 shows that the direct effect of an amenity shock is the mirror image of a product quality shock (-0.26 compared to 0.26). However, while variable markups and sectoral price indices dampen the passthrough of positive quality shocks, they amplify the passthrough of positive amenity shocks. This is because an increase in the value of amenities allows the firm to hire workers at a lower wage, effectively acting as a negative shock to production costs and expanding the size of the firm.

Given the above discussion on the role of markups, markdowns, bargaining, and sectoral price indices in determining the wage passthrough, a shock of a given size should have a different passthrough at different points in the wage premium distribution. Panel (B) of Table 5 shows that productivity and quality passthrough are substiantially smaller for firms in the 90th percentile, while the amenity passthrough is

larger. In comparison, firms at the 10^{th} and 50^{th} percentiles have a passthrough that is well-approximated by the direct effect only.

7.3 Decomposing the dispersion of firm wage premia

I now decompose the cross-sectional dispersion of firm wage premia and find that product quality is the main driver, closely followed by productivity and amenities.

To quantify the contributions of each source of firm heterogeneity to wage premium dispersion, I introduce *one* source at a time into the model. I then compare the implied variance of wage premia to that of the baseline economy, which exactly reproduces the wage premium distribution in the data. Table 6 shows the contribution of each source of heterogeneity.

Introducing heterogeneity in only product quality generates a larger wage premium dispersion than introducing only productivity or amenity heterogeneity. Product quality is thus a larger contributor to wage premium dispersion than productivity and amenities. Nevertheless, productivity and amenities play an important role in driving wage premia.

Panel (A)		Counterfactual	
	Vary only TFP	Vary only quality	Vary only amenity
$\frac{V(\phi^{\text{counterfactual}})}{V(\phi^{\text{baseline}})}$	152.8	160.7	40.6

Table 6: Importance of firm heterogeneity for firm wage premia.

This table reports the contribution of different sources of firm heterogeneity to the variance of firm wage premia. Each column shows the variance of wage premia when only one source of heterogeneity is present, compared to the baseline model in which all sources are active. Note that the variance of ϕ^{baseline} is the same as the variance of ϕ^{data} .

7.4 The impact of worker bargaining power and firm market power on wages and welfare

Section 6.2 discussed how firm product and labor market power distort social efficiency, and highlighted the role of worker bargaining power in mitigating these efficiency losses. In this section, I present a quantitative evaluation of these insights.

Wages and welfare. To compare the implications of variable markups/downs to constant markups/downs, I implement subsidy schemes that induce firms to choose the same markups/downs—without altering welfare through other channels—in an environment where workers have bargaining power. This extends the results in Edmond et al. (2015). The details of the derivations are in Appendix C.5.

I measure welfare in consumption-equivalent terms. Let household utility be U(C, H)in the baseline economy and $U(C^*, H^*)$ in a counterfactual economy. The consumptionequivalent welfare gain in percentages (\hat{c}) is measured as $U(C^*, H^*) = U((1 + \hat{c})C, H)$. Table 7 presents my findings.

Table 7: Effects of worker bargaining power and firm product and labor market power on wages and welfare.

		Counterfactuals									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)		
	$\mu_j = 1$	$\lambda_j = 1$	$\mu_j = 1, \\ \lambda_j = 1$	$\mu_j = \bar{\mu}$	$\lambda_j = \bar{\lambda}$	$\mu_j = \bar{\mu}, \\ \lambda_j = \bar{\lambda}$	$\kappa = 0.2$	$\kappa = 0.5$	$\kappa = 1.0$		
$\%\Delta V(\phi_j)$	98	23	112	97	24	112	5	26	49		
$\Delta \Phi$	97	83	302	39	10	55	20	97	225		
$\hat{c} \times 100$	38	9	46	24	2	27	3	10	13		

This table reports percentage changes in the variance of firm wage premia, the aggregate wage index, and aggregate welfare across a range of counterfactual simulations. Welfare gains \hat{c} are expressed in consumption-equivalent terms. Simulations with constant markups fix markups at $\mu = \frac{\sigma}{\sigma-1}$, while those with constant monopsony markdowns fix markdowns at $\lambda = \frac{\eta}{1+\eta}$.

To assess how far stronger worker bargaining power can go in improving welfare, I begin by evaluating the social planner's equilibrium, which eliminates both markups and monopsony markdowns. Column 3 shows that eliminating both markups and monopsony markdowns leads to a 112% increase in the dispersion of firm wage premia, a 302% increase in average wages, and a 46% welfare gain. Comparing columns 1 and 3 shows that over 80% of the welfare gain comes from removing markups.

Column 4 shows that removing markup dispersion alone—by setting constant CES markups ($\mu = \frac{\sigma}{\sigma-1}$)—delivers a 24% welfare gain and increases average wages and wage premia dispersion by 39% and 97%. This implies that about 63% of the welfare cost of markups stems from misallocation. In contrast, equalizing monopsony mark-downs has little effect on welfare (column 5), though a wide markdown *in levels* does reduce welfare significantly (column 2).

Can higher worker bargaining power (κ) offset these distortions? Column 8 shows that raising κ from 0.12 to 0.50 raises average wages by as much as removing markups (column 1), but with two key differences. First, it increases wage premia dispersion

much less, suggesting that worker power mainly offsets the *uniform tax* component of market power, not the misallocation from markup dispersion. Second, it yields a much smaller welfare gain (10% vs. 38%), because a higher κ corrects only distortions to labor demand, not those that directly affect capital or material demand.

For the same reason, giving workers full bargaining power does not restore the planner's equilibrium. This raises welfare by 13% (column 9), which is under one-third of the welfare gains from implementing the planner's equilibrium (column 3).

8 Conclusion

This paper develops a structural model to analyze how firm market power and worker collective bargaining power shape wages and welfare. A central insight is that the labor wedge between wages and the marginal revenue product of labor reflects not only monopsony power but also markups and bargaining power. As a result, markups can either amplify wage premia via rent-sharing or suppress wage dispersion by reducing labor demand. The model also shows that stronger worker bargaining power can improve welfare under firm market power. In addition, it highlights the challenges of separately estimating bargaining and monopsony power using existing approaches, and offers an alternative approach for doing so.

Quantifying the model yields several findings. First, the degree of product differentiation governs how productivity or quality shocks pass through to wages. Second, markup dispersion reduces wage premium dispersion, as its negative labor demand effect outweighs the rent-sharing channel. Third, while increasing worker bargaining power improves welfare, even full bargaining power cannot restore full efficiency, since it does not address the distortionary effect of markups on non-labor inputs.

Bibliography

- Abowd, J., F. Kramarz, and D. Margolis (1999). High Wage Workers and High Wage Firms. *Econometrica* 67(2), 251–333.
- Ackerberg, D., K. Frazer, and G. Caves (2015). Identification Properties of Recent Production Function Estimators. *Econometrica* 83(6), 2411–2451.

- Atkeson, A. and A. Burstein (2008). Pricing-to-Market, Trade Costs, and International Relative Prices. *American Economic Review* 98(5), 1998–2031.
- 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), 645–709.
- Berger, D., K. Herkenhoff, and S. Mongey (2022). Labor Market Power. American Economic Review 112(4), 1147–1193.
- Bhuller, M., K. Moene, M. Mogstad, and O. Vestad (2022). Facts and Fantasies about Wage Setting and Collective Bargaining. *Journal of Economic Perspectives* 36(4), 29–52.
- Bonhomme, S., K. Holzheu, T. Lamadon, E. Manresa, M. Mogstad, and B. Setzler (2023). How Much Should We Trust Estimates of Firm Effects and Worker Sorting. *Journal of Labor Economics*.
- Bonhomme, S., T. Lamadon, and E. Manresa (2019). A Distributional Framework for Matched Employer-Employee Data. *Econometrica* 87(3), 699–738.
- Burdett, K. and D. Mortensen (1998). Wage Differentials, Employer Size, and Unemployment. *International Economic Review* 39(2), 257–273.
- Cahuc, P., F. Postel-Vinay, and J. Robin (2006). Wage Bargaining with On-the-Job Search: Theory and Evidence. *Econometrica* 74(2), 323–364.
- Card, D., A. Cardoso, J. Heining, and P. Kline (2018). Firms and Labor Market Inequality: Evidence and Some Theory. *Journal of Labor Economics* 36(1), 13–70.
- Card, D., F. Devicienti, and A. Maida (2014). Rent-sharing, Holdup, and Wages: Evidence from Matched Panel Data. *Review of Economic Studies* 81(1), 84–111.
- Chetty, R. (2012). Bounds on Elasticities with Optimization Frictions: Synthesis of Micro and Macro: Evidence on labor Supply. *Econometrica* 80(3), 969–1018.
- De Loecker, J., J. Eeckhout, and G. Unger (2020). The Rise of Market Power and the Macroeconomic Implications. *Quarterly Journal of Economics* 135(2), 561–644.
- De Loecker, J. and P. Goldberg (2014). Firm Performance in a Global Market. *Annual Review of Economics 6*, 201–227.
- De Loecker, J., P. Goldberg, A. Khandelwal, and N. Pavcnik (2016). Prices, Markups, and Trade Reform. *Econometrica* 84(2), 445–510.
- De Loecker, J. and C. Syverson (2021). An Industrial Organization Perspective on Productivity. *Handbook of Industrial Organization IV*.

- De Loecker, J. and F. Warzynski (2012). Markups and Firm-Level Export Status. *American Economic Review* 102(6), 2437–2471.
- De Ridder, M., B. Grassi, and G. Morzenti (2021). The Hitchhiker's Guide to Markup Estimation. *Working paper*.
- Deb, S., J. Eeckhout, A. Patel, and L. Warren (2024). Market Power and Wage Inequality. *Econometrica* 92(3), 603–636.
- Di Addario, S., P. Kline, R. Saggio, and M. Solvsten (2020). It Ain't Where You're From, It's Where You're At: Hiring Origins, Firm Heterogeneity, and Wages. *Working paper*.
- Edmond, C., V. Midrigan, and D. Xu (2015). Competition, Markups, and the Gains from International Trade. *American Economic Review* 105(10), 3183–3221.
- Edmond, C., V. Midrigan, and D. Xu (2023). How Costly are Markups? *Journal of Political Economy* 131(7).
- Fieler, A., M. Eslava, and D. Xu (2018). Trade, Quality Upgrading, and Input Linkages: Theory and Evidence from Colombia. *American Economic Review 108*(1), 109–46.
- Gandhi, A., S. Navarro, and D. Rivers (2020). On the Identification of Gross Output Production Functions. *Journal of Political Economy* 128(8), 921–947.
- Garin, A. and F. Silverio (2024). How Responsive are Wages to Demand within the Firm? Evidence from Idiosyncratic Export Demand Shocks. *The Review of Economic Studies* 91(3), 1671–1710.
- Gouin-Bonenfant, E. (2022). Productivity Dispersion, Between-Firm Competition, and the Labor Share. *Econometrica* 90(6).
- Gourio, F. and L. Rudanko (2014). Customer Capital. *Review of Economic Studies* 81(3), 1102–1136.
- Hottman, C., S. Redding, and D. Weinstein (2016). Quantifying the Sources of Firm Heterogeneity. *Quarterly Journal of Economics* 131(3), 1291–1364.
- Jarosch, G., J. Nimczik, and I. Sorkin (2024). Granular Search, Market Structure, and Wages. *Review of Economic Studies (forthcoming)*.
- Jäger, S., B. Schoefer, S. Young, and J. Zweimüller (2020). Wages and the Value of Nonemployment. *Quarterly Journal of Economics* 135(4), 1905–1963.
- Klette, T. and Z. Griliches (1996). The Inconsistency of Common Scale Estimators When Output Prices are Unobserved and Endogenous. *Journal of Applied Econometrics* 11(4), 343–361.

- Kline, P., N. Petkova, H. Williams, and O. Zidar (2019). Who Profits from Patents? Evidence from Innovative Firms. *Quarterly Journal of Economics* 134(3), 1343–1404.
- Kline, P., R. Saggio, and M. Solvsten (2020). Leave-Out Estimation of Variance Components. *Econometrica* 88(5), 1859–1898.
- Kroft, K., Y. Luo, M. Mogstad, and B. Setzler (2024). Imperfect Competition and Rents in Labor and Product Markets: The Case of the Construction Industry. *Mimeo*.
- Krueger, A. and L. Summers (1988). Efficiency Wages and the Inter-Industry Wage Structure. *Econometrica* 56(2), 259–293.
- Kugler, M. and E. Verhoogen (2012). Prices, Plant Size, and Product Quality. *The Review of Economic Studies* 79(1), 307–339.
- Lamadon, T., M. Mogstad, and B. Setzler (2022). Imperfect Competition, Compensating Differentials, and Rent Sharing in the U.S. Labor Market. *American Economic Review* 112(1), 169–212.
- Mertens, M. (2020). Micro-Mechanisms Behind Declining Labor Shares: Market Power, Production Processes, and Global Competition. *Working paper*.
- Schmieder, J., T. Von Wachter, and J. Heining (2018). The Costs of Job Displacement over the Business Cycle and Its Sources: Evidence from Germany. *Working paper*.
- Sokolova, A. and T. Sorensen (2020). Monopsony in Labor Markets: A Meta-Analysis. *ILR Review* 74(1), 27–55.
- Sorkin, I. (2018). Ranking Firms Using Revealed Preference. *Quarterly Journal of Economics* 133(3), 1331–1393.
- Stansbury, A. and L. Summers (2020). The Declining Worker Power Hypothesis. *Brookings Papers on Economic Activity*.
- Taschereau-Dumouchel, M. (2020). The Union Threat. *Review of Economic Studies* 87(6), 2859–2892.
- Van Reenen, J. (1996). The Creation and Capture of Rents: Wages and Innovation in a Panel of U.K. Companies. *Quarterly Journal of Economics* 111(1), 195–226.
- Webber, D. (2015). Firm Market Power and the U.S. Earnings Distribution. *Labor Economics* 35, 123–134.
- Yeh, C., C. Macaluso, and B. Hershbein (2022). Monopsony in the US Labor Market. *American Economic Review* 112(7), 2099–2138.

Appendix (for online publication)

Understanding High-Wage Firms: Monopoly, Monopsony, and Bargaining Power

Horng Chern Wong Department of Economics, Stockholm University Email: horngchernwong@gmail.com

A Appendix: Institutions, Data, and Measurement

A.1 Wage determination in France

Wages in France are mainly determined at three levels of aggregation—the national level, the industry level, and the firm level. At the national level, the French government sets the national minimum wage. At the industry level, labor unions and employers' organizations negotiate industry wage floors. At the firm level, wage bargaining occurs between an individual firm and labor union representatives, representing the collective group of employees at the firm. This allows firms to depart upwards of the national minimum wage or industry wage floor.

The national minimum wage as of 2016 is 9.67 Euros per hour. In the same year, approximately 15% of workers in the French economy are minimum wage workers, defined as workers earning at, or at most 5% above, the minimum wage. In the manufacturing sector, on which the analysis in the paper is based, approximately 10% of workers are minimum wage workers.

Collective wage bargaining between individual firms and their employees is prevalent. The 1982 Auroux Laws require firms in which a labor union representative is present to bargain over wages with the union annually. Among firms with 50 employees or more, the presence of at least one union representative is a binding legal requirement.²³ Even among firms with fewer than 50 employees, the presence of union representatives is common. According to the French Ministry of Labor, among firms with

²³See Garicano, Van Reenen, and Lelarge (2016) for further details about the specific restrictions faced by firms with at least 50 employees.

20 to 49 employees, 34% had at least one labor union representative in 2010 (Naouas and Romans, 2014). Among firms with 11 to 19 employees, the corresponding number is 22%.

The vast majority of union-employer bargaining occurred at the firm level, rather than at the establishment level. Only 9% of multi-establishment firms negotiated wages with their employees at the establishment level. Among workers employed by firms with at least 20 employees, 70% of them are covered by firm level collective bargaining agreements. These collective bargaining agreements extend to all workers within the firm, regardless of whether the worker holds a union membership (Fougere, Gautier, and Roux, 2016).

A.2 Data and measurement

To estimate production functions using FARE-EAP-DADS (2009-2016) firm balance sheet, output price, and matched employer-employee data from France, I measure the key variables in the following way:

- ► Sales revenue (*PY*): measured by the variable CATOTAL in FICUS, and REDI_R310 in FARE.
- ► Efficiency units of labor (*H* = *ĒL*): the DADS provides the number of hours worked for each worker under NBHEUR, which enables the researcher to measure total hours (*L*) at a given firm. The average efficiency of workers (*Ē*) is then measured as the difference between the unconditional mean wage and the firm wage premium, according to the theory.
- ► Capital (*K*): measured as total fixed physical assets under variable names IM-MOCOR in FICUS, and IMMO_CORP in FARE.
- Materials (*M*): the French balance sheet data provides a breakdown of material inputs into three components materials purchased to be used as inputs in production (ACHAMPR in FICUS, REDI_212 in FARE), goods purchased to be resold (ACHAMAR in FICUS, REDI_210 in FARE), and purchase of services (details provided next). I correct for changes in inventory for materials to be used in production (using VARSTMP in FICUS, REDI_213 in FARE) and for goods pur-

chased to be resold (VARSTMA in FICUS, REDI_211 in FARE). I measure *M* as the sum of these variables, except services.

- Services (*O*): measured as AUTACHA in FICUS, and REDI_214 in FARE. These variables include the costs of outsourcing and advertising.
- ► Hourly wages (*W*): measured by dividing BRUT by NBHEUR in DADS.
- Output prices (P): PRODFRA defines the ten-digit product codes, C_UNITE_VAR gives the quantity and revenue indicator, and VAL_REF gives the values in quantity and revenue terms. These variables are obtained from EAP.
- ▶ Market shares: measured within 5-digit sectors.

A.3 Estimation sample

I restrict firm-level observations from the FARE data to manufacturing firms whose output prices are observed in the EAP. I include only firms with at least 5 employees. I harmonize all industry codes to the latest available version (Nomenclature d'activités Française, NAF rév. 2) and exclude 2-digit sectors with fewer than 300 observations.

For both of the DADS datasets, I focus on workers aged 16 to 65, who hold either a part-time or full-time principal job (jobs in which workers are paid for at least 30 days of work and at least 120 hours of work that year). I use only years 2009-2016, since INSEE reports that wages are recorded with errors in 2017 and 2018. I keep workers in almost all 1-digit occupational categories, except farm workers. The included occupational categories are top management, senior executives and technical professions, middle management, non-supervisory white-collar workers, and blue-collar workers. Occupation codes are harmonized and updated to the latest version (PCE-ESE 2003). Workers whose wages fall outside 3 standard deviations of the mean are excluded.

Firm wage premia in the wage regression are only identified for the sets of firms connected by worker mobility. I focus on the largest connected set of firms. In practice, due to the clustering of firms into groups using the DADS-Postes, my analysis pertains to the largest connected set of firm-groups, of which very few firms are not a part. After clustering firms into groups, I link the DADS-Postes and DADS-Panel via the firm identifier to allocate each firm-year observation a firm-group identifier and construct the estimation sample for firm wage premia. I estimate firm wage premia on this sample. Figure 22 in Appendix D shows that larger firms pay higher wage premia.

After estimating firm wage premia, I collapse the dataset to the firm level and link it to the FARE-EAP balance sheet and output price data to construct the estimation sample. I implement the production function estimation routine on this sample. The summary statistics in Table 8 in Appendix A.4 show that there is significant price dispersion within narrowly defined sectors. Figure 23 in Appendix D further shows that this price dispersion is positively correlated with firm size.

A.4 Summary statistics

	Sales (′000€)	Employment	Sales per hour (′000€)	Prices	Prices (2-digit)	Prices (5-digit)
10 th percentile	2,561	8	0.08	-1.50	-1.21	-0.87
25 th percentile	3,794	13	0.10	-0.30	-0.34	-0.31
Median	5,853	27	0.15	0.26	0.14	0.07
Mean	24,237	90	0.19	0.00	0.00	0.00
75 th percentile	12,355	59	0.22	0.49	0.46	0.37
90 th percentile	34,780	167	0.33	0.84	0.91	0.82
Std. deviation	148,810	369	0.40	0.97	0.89	0.81

Table 8: Summary statistics of observed firm characteristics.

This table reports the summary statistics for firm-level sales, employment, sales per hour, and output prices. Sales and sales per hour are denoted in 2009 prices. Output prices are in logs, normalized to zero at the mean, and winsorized by 5% on either side of the distribution. The last two columns, with '2-digit' and '5-digit' in parentheses, display variation in log output prices within 2-digit and 5-digit sector classifications.

	BI	LM	A	KM	KSS		
	MN	All	MN	All	MN	All	
$\frac{Var(\phi)}{Var(w)}$	3.7%	5.2%	10.6%	15.8%	4.6%	7.2%	
$Var(\phi)$	0.008	0.011	0.016	0.024	0.006	0.011	
90-10 ratio	1.23	1.30	1.35	1.43	-	-	
75-25 ratio	1.12	1.15	1.16	1.20	-	-	
90-50 ratio	1.10	1.14	1.15	1.20	-	-	
50-10 ratio	1.12	1.14	1.17	1.19	-	-	
# firms	14,599	271,272	9,669	167,078	9,669	167,078	
# firm-groups	399	3,276	9,669	167,078	9,669	167,078	

Table 9: Dispersion of firm wage premia.

This table reports the dispersion of firm wage premia in 2016. Columns labeled 'MN' are estimates for manufacturing firms while columns labeled 'All' include all private sector firms in my sample. Columns labeled 'BLM' apply the Bonhomme et al. (2019) clustering approach, columns labeled 'AKM' apply Abowd et al. (1999), and columns under 'KSS' apply the Kline et al. (2020) leave-out approach.

Sector	# Observations	Sales share	Employment share	Average ϕ	Price-cost markups	Labor wedges
Textile	4,902	1.5%	2.1%	2.88	1.36	0.69
Apparel	4,060	1.7%	2.5%	2.85	1.40	0.71
Leather	1,695	1.4%	1.7%	2.80	1.31	0.86
Wood products (except furniture)	6,161	1.6%	2.3%	2.85	1.19	0.54
Paper and publishing	4,094	2.7%	2.8%	2.89	1.27	0.57
Recorded media	7,573	1.3%	2.3%	2.90	1.67	0.79
Chemicals	6,417	16.7%	9.5%	2.96	1.24	0.47
Pharmaceutical	381	4.7%	2.0%	3.01	1.10	0.20
Rubber & plastics	12,662	9.1%	11.6%	2.91	1.29	0.61
Non-metallic minerals	8,194	5.9%	6.5%	2.90	1.27	0.54
Basic metals	3,408	7.4%	5.5%	2.94	1.25	0.55
Fabricated metals (except machinery)	21,145	6.9%	10.0%	2.91	1.37	0.62
Computers, electronic, & optical	4,123	5.5%	6.2%	2.95	1.28	0.62
Electrical equipment	5,463	7.2%	7.4%	2.92	1.23	0.58
Machinery & equipment	10,204	9.7%	9.8%	2.94	1.11	0.44
Motor vehicles	3,927	7.3%	7.0%	2.94	1.14	0.55
Other transport equipment	668	4.4%	3.3%	2.92	1.05	0.41
Furniture	7,191	1.8%	3.0%	2.88	1.42	0.72
Other manufacturing	2,984	1.7%	2.3%	2.91	1.36	0.70
Repair & installation of machinery	3,669	1.5%	2.2%	2.89	1.37	0.58
Total	118,921	100%	100%	-	-	-

Table 10: Summary statistics for 2-digit French manufacturing sectors.

This table reports the summary statistics for manufacturing sectors in my sample (2009-2016). The last two columns report the average price-cost markup and labor wedges in each sector.

B Appendix: Estimation

B.1 K-means clustering of firms into groups

Specifically, let $g(j) \in \{1, 2, ..., G\}$ denote the cluster of firm j, and G the total number of clusters. The k-means algorithm finds the partition of firms such that the following objective function is minimized:

$$\min_{g(1),\dots,g(J),H(1),\dots,H(G)} \sum_{j=1}^{J} N_j \int \left(\hat{F}_j(\ln W_{ij}) - H_{g(j)}(\ln W_{ij}) \right)^2 d\gamma(\ln W_{ij})$$

where H(g) denotes the firm-group level cumulative distribution function for log wages at group g, \hat{F}_j is the empirical CDF of log wages at firm j, and N_j is the employment size of firm j. The total number of groups G is the choice of the researcher. I choose sector-specific G such that the variance of log wages between firm-groups captures at least 95% of the unconditional between-firm variance. This choice is motivated by the following consideration: having a coarse classification of firms into fewer groups leads to many more workers who switch between firm-groups, which substantially improves the precision of firm wage premium estimates. However, this comes at the cost of potentially averaging away considerable amounts of multidimensional firm heterogeneity within firm-groups.

B.2 AKM restrictions: Exogenous mobility and log-additivity

AKM regressions rely on the assumption that worker mobility is as good as random conditional on observed worker characteristics, worker fixed effects, and firm fixed effects. Formally, $E(v_{it}|\chi_{it}, \iota_i, \phi_{g(j(i,t),t)t}) = 0$. This assumption rules out worker mobility based on wage realizations due to the residual component of wages. In addition, AKM regressions impose log additivity of the worker and firm components of wages. If these assumptions are reasonable approximations, then one should observe systematic worker mobility up and down the firm wage quartiles. Moreover, workers should experience approximately symmetric wage changes as they move along the firm wage quartiles, given the log additive regression specification. On the other hand, in structural models of worker-firm sorting based on comparative advantage (Eeckhout and

Kircher, 2011), worker mobility is based on the match-specific component of wages, which is captured by the residual component of wages in the AKM regression. In this class of models the AKM regression is misspecified in the sense that the wage gains depend on value of the particular worker-firm match, for example, if highly skilled workers have a comparative advantage in high productivity firms. In the event-study exercise show in Figure 3, I compare the changes in mean log wages for workers who move between firms in different quartiles of coworker pay, following Card et al. (2018). Figure 3 shows that workers who move up firm quartiles experience a wage gain similar in magnitude to the wage loss of workers who move down firm quartiles.

An alternative way to assess the AKM regression specification is to compare the changes in residual wages to changes in firm effects, following Sorkin (2018). This is similar to the above method. I run the following regression among all employer-to-employer transitions:

$$w_{it}^{r} - w_{it-1}^{r} = \alpha_{0} + \alpha_{1} \left(\phi_{g(j(i,t))} - \phi_{g(j(i,t-1))} \right) + \epsilon_{it} \,\,\forall \,(i,t), \,\,g(j(i,t)) \neq g(j(i,t-1))$$

where $w_{it}^r = w_{it} - x_{it}'\hat{\beta}$ denotes residualized wages and $\phi_{g(j(i,t))}$ are the firm-group fixed effects. If the AKM regression is not mis-specified, the estimated coefficient $\hat{\alpha}_1$ will equal 1. I find $\hat{\alpha}_1 = 0.902$, with a standard error of 0.012. To see this visually, Figure 4 plots the changes in residual wages and the changes in firm fixed effects in 100 bins of changes in firm fixed effects. In models of assortative matching based on comparative advantage (Lopes de Melo, 2018), worker mobility is strongly driven the residual component of the AKM regression, implying that AKM regressions are misspecified. As Sorkin (2018) shows, these models predict that worker mobility entails a wage gain, regardless of the direction of mobility in terms of the estimated firm effects, as workers move to firms at which they have a comparative advantage: there is a V-shape around zero changes in firm effects. The patterns of wage changes upon changes in firm fixed effects shown in Figure 4 do not resemble a V-shape around zero.

Another way to assess the log additivity assumption is to group worker and firm fixed effects into 10 deciles each, generating 100 worker-firm fixed effect deciles, then plot the mean estimated residuals within each worker-firm fixed effect decile. If the firm wage premium depends strongly on the worker's unobserved ability type, log additivity would be severely violated, and one should observe that the estimated residuals systematically varies across worker-firm fixed effect deciles. Figure 5 shows that the mean estimated residuals are approximately zero across worker-firm fixed effect deciles, with the exception of the very top deciles of high-wage workers who are employed at low-wage firms at the very bottom deciles.

As a further robustness check, I follow Bonhomme et al. (2019) and run the BLM regression with worker-firm interactions, but with only 20 firm groups and 6 worker groups to maintain computational tractability.



Figure 3: Worker mobility and wage changes by quartiles of coworker effects.



Figure 4: Wage changes from worker mobility by deciles of changes in firm premia.



Figure 5: Mean estimated residuals by worker-firm deciles (2016)

B.3 Unobserved prices in production function estimation

When estimating production functions, a common challenge researchers face is that output and input prices are not observed. As discussed in detail by Klette and Griliches (1996) and De Loecker and Goldberg (2014), these unobserved prices may lead to biased estimates of output elasticities as well as firm productivity.²⁴ The French microdata contains information about the prices and quantities of goods sold by manufacturing firms, making it possible to address unobserved output prices. These output price data also make it possible to address unobserved input price heterogeneity, as De Loecker et al. (2016) show. In this section, I explain how I address these challenges. In Appendix B.4, I find that addressing neither of the biases that stem from unobserved output and input prices turns out to produce similar estimates of output elasticities as estimations that attempt to address both.

Unobserved output prices. Unobserved output prices may introduce an *output price bias* into the estimation of production function parameters because output prices are correlated with input choices (Klette and Griliches, 1996).²⁵ In practice, the production function that is frequently estimated takes the form (abstracting from input price heterogeneity for now):

$$p_{jt} + y_{jt} = f_s(k_{jt}, h_{jt}, m_{jt}; \beta) + p_{jt} + \omega_{jt}$$

where $p_{jt} + \omega_{jt}$ is revenue TFP (TFPR) and ω_{jt} is quantity TFP (TFPQ). Intuitively, all else equal firms that set higher prices tend to sell less output, reducing the demand for inputs. This potential negative correlation between output prices and input demand introduces a downward bias in the estimated output elasticities. In practice, to address unobserved output prices, researchers have obtained price data (the approach taken in this paper) and/or added more structure on the product demand-side (see De Loecker and Goldberg (2014) for a detailed discussion of the relevant tradeoffs).

However, unobserved output prices do not always bias the estimated output elasticities. As De Loecker and Warzynski (2012) explain, if price variation reflects Hicksneutral productivity, it is absorbed by the productivity control function; only demand-

²⁴Although the production approach to markup estimation requires estimates of output elasticities, this approach does not rely on any one particular method of estimating those elasticities.

²⁵When output prices are observed, they are typically specific to certain industries, such as beer brewing (De Loecker and Scott, 2024).

or quality-driven price variation uncorrelated with TFPQ introduces bias. Despite this potential bias, revenues can be more appropriate than quantities when comparing differentiated products, provided that quality-driven price variation is controlled for. De Loecker et al. (2016) explicitly incorporate quality differences into the production function and develop a framework to address them (more on this below).

Nevertheless, the output price data for French manufacturing firms is useful for several reasons. First, the data makes it possible to compute a firm-level price index to directly address the output price bias (see Section 4). Second, the data helps account for quality differences uncorrelated with TFPQ and address the associated input price heterogeneity (De Loecker et al., 2016). Third, the data makes it possible to measure TFPQ and quality to quantify their separate contributions to wage dispersion—a goal of the quantitative analysis in Section 7.

I measure firm-level prices p_{jt} using French administrative data on firm-productyear level prices for manufacturing firms. I then compute firm-year level output y_{jt} as revenue divided by my measure of firm-level prices. As De Loecker and Goldberg (2014) explain in detail, although the availability of output prices may help alleviate the biases from unobserved prices, incorporating them into production function estimation also brings other challenges. In particular, the output price data are often at the firm-product level; properly aggregating them into firm-level price indices requires a theory of multi-product firms that goes beyond the scope of this paper. In Section 4, I provide further detail on a measure of standardized firm-level output prices that I implement, proposed by De Ridder et al. (2021). Appendix B.4 shows that the markup and labor wedge estimates are not materially affected by whether output is measured in revenues or (a measure of) quantities.

Unobserved input prices. Researchers rarely observe firms' input prices directly, so deflated input expenditures are typically used in place of quantities. This potentially introduces an *input price bias* in the estimated production function parameters (De Loecker and Goldberg, 2014). Conditional on observing firms' output prices, the commonly estimated production function becomes:

$$y_{jt} = f_s(\tilde{k}_{jt}, \tilde{h}_{jt}, \tilde{m}_{jt}; \beta) + B(\mathbf{p}_{\mathbf{x}, jt}, \tilde{\mathbf{x}}_{jt}; \beta, \zeta) + \omega_{jt}$$

where $\tilde{\mathbf{x}}_{jt} = {\{\tilde{k}_{jt}, \tilde{h}_{jt}, \tilde{m}_{jt}\}}$ represents the set of input expenditures (denoted with a

tilde) and $\mathbf{p}_{\mathbf{x},jt} = \{p_{k,jt}, \phi_{jt}, p_{m,jt}\}$ the set of input prices. The function B(.) captures the influence of unobserved input prices, with its functional form depending on the functional form of the production function f(.). The parameters ζ must then be estimated to account for unobserved input prices.

Relative to most datasets, the French administrative data includes hours and wages at the worker level, allowing me to measure effective labor h_{jt} . However, capital and material quantities are not observed. The production function I estimate becomes:

$$y_{jt} = f_s(k_{jt}, h_{jt}, \tilde{m}_{jt}; \beta) + B_s(\mathbf{p}_{\mathbf{x}, jt}, \tilde{\mathbf{x}}_{jt}, h_{jt}; \beta, \zeta) + \omega_{jt}$$

where the set of input expenditures is now $\tilde{\mathbf{x}}_{jt} = {\tilde{k}_{jt}, \tilde{m}_{jt}}$ and the set of unobserved input prices is $\mathbf{p}_{\mathbf{x},jt} = {p_{k,jt}, p_{m,jt}}$. Because the prices of material and capital inputs are unobserved, as is the case in most existing datasets, I work under the standard assumption that firms are pricetakers in these input markets (De Loecker and Goldberg, 2014). To the extent that firms in different sectors and locations face different material and capital input prices, I control for sector and location fixed effects in the production function estimation routine.²⁶

However, input prices may differ across firms due to differences in input quality. As De Loecker et al. (2016) show, under a large class of consumer demand functions used in International Trade, Industrial Organization, and Macroeconomics, output prices increase monotonically with output quality, which increases monotonically with input quality. Under the assumption that higher quality inputs come with higher input prices, one can then build a control function for unobserved input prices using output prices. Specifically, let input prices $\mathbf{p}_{\mathbf{x},jt} = \mathbf{p}_{\mathbf{x}}(d_{jt}, \mathbf{G}_j)$ depend on output quality d_{jt} and fixed sector-location characteristics \mathbf{G}_j . De Loecker et al. (2016) show that the control function for input prices $\mathbf{p}_{\mathbf{x},jt} = \mathbf{p}_{\mathbf{x}}(p_{jt}, \mathbf{Z}_{jt})$ can be written as a function of output prices p_{jt} and a vector \mathbf{Z}_{jt} containing sector-location fixed effects \mathbf{G}_j .

Relative to De Loecker et al. (2016), I allow labor markets to be imperfectly competitive. This implies that monopsony markdowns will affect output prices through marginal costs. Therefore, for output prices to be valid proxies for input prices, the input price control function needs to account for monopsony markdowns. I do so by including firm wage premia in \mathbf{Z}_{jt} in the input price control function. The function

²⁶This can be the case, for example, due to differences in market access by location.

B(.) can then be written as:

$$B_s((p_{jt}, \mathbf{Z}_{jt}) \times \{1, \tilde{\mathbf{x}}_{jt}, h_{jt}\}; \beta, \zeta)$$

which is a function of output prices p_{jt} and the vector of controls \mathbf{Z}_{jt} , and their interactions with input expenditures $\tilde{\mathbf{x}}_{jt}$ and effective labor h_{jt} . Since input expenditures $\tilde{\mathbf{x}}_{jt}$ only enter the function B(.) as interaction terms with output prices and other controls \mathbf{Z}_{jt} , the production function parameters β are identified. This identification insight from De Loecker et al. (2016) does not hinge on functional form assumptions for f(.).²⁷

²⁷I refer interested readers to De Loecker et al. (2016) for derivations of this identification result. The reason that input expenditures do not enter B(.) as lone variables is that the control function for input prices is built only from the consumer demand side.

B.4 On the role of data on hours and prices

The French matched employer-employee data (DADS) and production survey of manufacturing firms (EAP) offer three key pieces of information: worker-level hours worked, worker mobility between employers, and output prices. Data on hours worked allows me to measure labor inputs as total hours. Data on workers' employer switches allows me to measure the total effective hours at a firm as the difference between the firmlevel average hourly wage and firm wage premium (see Section 2). Data on output prices allows me to measure output in quantities instead of revenues, to address unobserved input price heterogeneity stemming from input quality differences (De Loecker et al., 2016), as well as to include controls for markups and output prices in the productivity control function. I compare the labor wedges and markups that I obtain under different specifications or measurements of labor inputs and output and arrive at two main findings.

1. Role of hours data: I find that measuring labor as total hours (or effective hours) instead of employment results in a lower *level* and smaller *dispersion* of labor wedges and markups. This finding may be useful to keep in mind when interpreting production-based labor wedge and markup measures, particularly in light of an important insight of recent macroeconomic models of labor and product market power: the level of labor wedges or markups act as a uniform proportional tax on firms, while the dispersion leads to resource misallocation across firms (Berger et al., 2022; Edmond et al., 2023).

Rows (1), (2), and (3) of Table 11 present the results when I estimate a basic specification that *does not* account for output price bias, input price bias, or markups in the productivity control function—only the measure of labor inputs differs. The mean and median labor wedge are lower when labor inputs are measured as total hours or total effective hours. The interquartile range and 90-10 difference for both labor wedges and markups are also smaller when hours data is used. Rows (7), (8), and (9) show that the same findings emerge in my baseline specification, which addresses output price bias, input price bias, and the presence of markups and output prices in the productivity control function.

To see how failing to account for variation in labor hours across firms could bias labor wedge and markup estimates, consider the following example. For brevity, I drop all firm and time subscripts. Let l, h, and \bar{h} denote employment, total hours, and

average hours in logs at a given firm. Suppose the production function consists of only labor and materials (m):

$$y = f(m,h) + \omega = f(m,l+\bar{h}) + \omega$$

I now discuss two sources of bias when data on hours is not available.





Note: This figure presents how log average hours vary across firm size deciles, controlling for 2-digit sector×year fixed effects. Vertical bars are 95% confidence intervals.

OMITTED VARIABLE BIAS. This equation shows that measuring labor inputs as employment instead of total hours leads to average hours (\bar{h}) being an omitted variable. If workers at larger firms work longer hours, then the production function parameter estimates (and returns to scale) will be biased upwards. Figure 6 shows that, within two-digit manufacturing sectors, average hours are higher at larger firms. In this case, material output elasticities will be upward-biased, leading to upward-biased markups. However, labor output elasticities will also be upward-biased, so it is ambiguous in which direction measured labor wedges will be biased. When I use employment instead of total effective hours, I find a 0.02 increase in the median returns to scale.

MISMEASUREMENT. When output elasticities depend on labor inputs, mismeasurement of labor inputs will also bias the output elasticities, even if production function parameters are known. The bias in measured output elasticities then lead to biased labor wedges and markups. Consider the following more specific example. Let f(m,h) be translog with constant returns to scale:

$$y = \beta m + (1 - \beta)h - \frac{\sigma_{mh}}{2}(m - h)^2 + \omega$$

where σ_{mh} is the elasticity of substitution between materials and labor. Materials and labor are substitutes when $\sigma_{mh} < 0$, and complements when $\sigma_{mh} > 0$. Suppose the production function parameters (β , σ_{mh}), TFPQ (ω), materials (m), and employment (l) are known, but not average hours \bar{h} . Let measured output elasticities be denoted with a hat. Then, the relationship between actual and measured output elasticities are:

Materials:
$$\hat{\alpha}_m = \alpha_m + \sigma_{mh}\bar{h}$$

Labor: $\hat{\alpha}_h = \alpha_h - \sigma_{mh}\bar{h}$

where $\sigma_{mh}\bar{h}$ is the source of the mismeasurement. In this case, the returns to scale is not affected. However, the sign of the bias on materials and labor output elasticities go in opposite directions and depends on the elasticity of substitution between materials and labor. Suppose materials and labor are complements ($\sigma_{mh} > 0$). Then, measured material output elasticities will be upward-biased, while measured labor output elasticities will be downward-biased. In turn, this implies that both measured markups ($\hat{\mu}$) and labor wedges ($\hat{\Lambda}$) will be biased upwards—recall that $\hat{\mu} = \hat{\alpha}_m \frac{PY}{P_mM}$ and $\hat{\Lambda} = \frac{\hat{\alpha}_m}{\hat{\alpha}_h} \frac{WL}{P_mM}$. These biases are consistent with the results from comparing rows (1), (2), and (3), as well as the comparison of rows (7), (8), and (9).

2. Role of output price data: I find that output and input price biases appear to partially offset each other. Specifications that address neither of the biases that stem from unobserved output nor input prices deliver broadly similar labor wedges and markups as specifications that address both. However, addressing only output price bias without addressing input price bias leads to biased labor wedges and markups.

I start with the comparison between row (3) and row (9). In row (3), I implement a "basic" version of production function estimation that *does not* address potential output and input price bias, or the presence of markups and output prices in the productivity control function. In contrast, the baseline implementation in row (9) addresses all three potential sources of bias. In both rows (3) and (9), I measure labor inputs as total effective hours. The measured labor wedges and markups are slightly higher in levels and dispersion with the basic implementation compared to my baseline implementation.

Comparing row (3) with row (4), where the only difference is that row (4) measures output in quantity terms, leads to substantially different levels and dispersion of markups and labor wedges. Correcting for output price bias appears to lead to lower and more dispersed markups, and higher and more dispersed labor wedges. In particular, correcting only for output price bias increases the median labor wedge from 0.64 to 0.77. Comparing row (4) with row (5) shows that controlling for output prices and markup proxies (market shares, firm age, and export status) in the productivity control function reduces the measured labor wedges. Comparing row (4) with row (6) shows that controlling for unobserved input price variation using output prices, following De Loecker et al. (2016), returns similar labor wedges and markups as the "basic" and baseline implementations in rows (3) and (9).

Why might addressing the output price bias only, but not the input price bias, lead to significantly lower markups and higher labor wedges? In what follows, I show that this could be the case when these two sources of bias offset each other, a point discussed in detail in De Loecker and Goldberg (2014) and De Loecker and Syverson (2021). Consider again the production function:

$$y = f(m, h) + \omega$$
$$\Rightarrow \tilde{y} = f(\tilde{m} - p_m, h) + \omega + p$$

where \tilde{y} and \tilde{m} represent log sales and log material expenditure. Let both output prices p = p(d) and material prices $p_m = p_m(d)$ be a function of output quality d.

OUTPUT PRICE BIAS. Suppose productivity, material quantities, and labor are observed, but not output prices p. Then, the production function estimates are *upward-biased* if output quality d is not perfectly correlated with productivity ω . This is because, all else equal, firms with higher output quality attract higher demand, charge higher prices, and demand more inputs, resulting in a positive correlation between input choices and unobserved output prices.

INPUT PRICE BIAS. Suppose productivity, output quantities, and labor are observed, but not material prices p_m . Then, the production function estimates are *downward-biased* if output quality *d* is not perfectly correlated with productivity ω . This is because, all else equal, firms with higher output quality attract higher demand and purchase higher-quality material inputs. This generates a positive correlation between

material choices and unobserved material prices.

Indeed, in the basic specification that does not address any of the potential sources of bias (row (3)), I find a median returns to scale of 1.01, which reduces to 0.97 when I address only the output price bias (row (4)). This suggests the presence of input price bias. Although there is no reason ex-ante to expect the output and input price biases to offset each other, the comparison between rows (3), (4), (5), (6), and (9) appear to support this interpretation.

				Markups			Labor wedges				
	Specification	<i>h</i> measure	<i>y</i> measure	Mean	Median	75-25	90-10	Mean	Median	75-25	90-10
(1)	Basic	Emp.	Sales	1.46	1.41	0.54	1.14	0.71	0.71	0.34	0.70
(2)	Basic	Hours	Sales	1.43	1.39	0.52	1.10	0.67	0.68	0.31	0.64
(3)	Basic	E. hours	Sales	1.39	1.37	0.52	1.05	0.62	0.64	0.29	0.60
(4)	Basic	E. hours	Quantity	1.39	1.32	0.54	1.19	0.80	0.77	0.55	1.19
(5)	Markup+price in control function	E. hours	Quantity	1.34	1.32	0.73	1.98	0.72	0.66	0.76	1.80
(6)	Input price controls only	E. hours	Quantity	1.28	1.32	0.48	1.05	0.55	0.61	0.29	0.65
(7)	Baseline	Emp.	Quantity	1.40	1.37	0.50	1.07	0.68	0.68	0.32	0.65
(8)	Baseline	Hours	Quantity	1.37	1.35	0.49	1.03	0.64	0.65	0.29	0.59
(9)	Baseline	E. hours	Quantity	1.33	1.33	0.47	0.97	0.60	0.62	0.26	0.51

Table 11: Measures of markups and labor wedges with and without data on hours and prices.

This table reports the measured markups and labor wedges when production functions are estimated with and without data on hours and/or output prices. All markup and labor wedge measures are trimmed by 1% on either side of their distributions. The 'baseline' specification refers to the specification estimated in Section 3, which includes output prices and controls for markups in the control function, and uses output prices to address unobserved quality-driven input price variation (De Loecker et al., 2016). The 'basic' specification neither controls for markups and output prices in the control function, nor addresses unobserved input quality differences. In row (5), the specification 'Markup+price in control function' includes output prices and controls for markups (i.e., market shares, firm age, export status) in the control function, but does not address input price bias. In row (6), the specification 'Input price controls only' addresses both input and output price biases, but does not address the presence of markups in the control function, leading to violations of the scalar unobservable assumption. The columns '*h* measure' and '*y* measure' show the measure of labor and output used in the estimated production function. 'E. hours' means effective hours.

B.5 On the identification of worker bargaining power

In Section 3.3, I propose a method for estimating workers' bargaining power (κ) based on the labor wedge equation (3), which specifies the theory-implied relationship between labor wedges and product market rents (henceforth, "markups"). One key challenge in identifying κ using this equation is the need to address unobserved amenities (to the extent that they determine monopsony markdowns). In that section, I propose a control function approach to address this challenge, using data on wages and employment to control for variation in amenities.

A natural question is: once wages and employment are included as controls, is there any variation left in markups to identify workers' bargaining power? In this section, I provide an example to illustrate that if firm productivity (Ω_j) and product quality (D_j) are not perfectly correlated—both between firms as well as within firms over time—then, in theory, there will be residual variation in markups to identify κ .

For this purpose, I now revisit the model presented in Section 2, but with a minor simplification: firms' production functions now depend only on employment (in efficiency units), $Y_j = \Omega_j H_j$. Then, solving the collective bargaining problem posed in Section 2 yields the following labor wedge equation: $\Lambda_j = \kappa \mu_j + (1 - \kappa)\lambda_j$, where markups $\mu_j = \mu(Y_j, D_j)$ depend on firm size (output) and product quality, and monopsony markdowns $\lambda_j = \lambda(H_j, A_j)$ depend on firm size (employment) and amenities. Since the model features firm heterogeneity in productivity (Ω_j), amenities (A_j), and product quality (D_j), the solution for firms' choice of employment and wages can be written as:

$$H_j = H(\Omega_j, A_j, D_j)$$
 and $\Phi_j = \Phi(\Omega_j, A_j, D_j)$

which, given the monotonicity assumptions made in Section 2, can be inverted and combined to yield:

$$\Omega_j = \Omega(H_j, \Phi_j, D_j) \quad \text{and} \quad A_j = A(H_j, \Phi_j, D_j) \tag{B.5-1}$$

These equations show the implied set of (Ω_j, A_j) for a given set of (H_j, Φ_j, D_j) .

To see that there there will be residual variation in markups to identify κ , consider two firms, j and j', with the same observed wage and employment: $\Phi_j = \Phi_{j'} = \bar{\Phi}$ and $H_j = H_{j'} = \bar{H}$. Substituting the equations (B.5-1) into the labor wedge equation yields:

$$\Lambda_{i} = \kappa \mu(\Omega(\bar{H}, \bar{\Phi}, D_{i})\bar{H}, D_{i}) + (1 - \kappa)\lambda(\bar{H}, \bar{\Phi})$$

Therefore, as long as firm productivity and product quality are not perfectly correlated, two firms with the same observed employment and wages can charge different markups—there is residual variation in markups for identifying κ .

Nevertheless, in practice it could be that, once firm size and wages are controlled for, there is little variation left in markups to identify κ . Reassuringly, Table 2 shows that the bargaining power estimates are stable across specifications that include or omit labor and wages.

B.6 The choice of production functions

B.6.1 Translog vs Cobb-Douglas

In Section 3, I estimate translog production functions to obtain firm-specific output elasticities, which are used to measure markups and labor wedges. The Cobb-Douglas production function is a simpler alternative that allows one to document how markups and labor wedges vary across firms without having to estimate production functions. The advantage is that one can sidestep the challenges of estimating gross output production functions (Gandhi et al., 2020). However, this requires assuming that output elasticities do not vary across firms.

In this section, I provide suggestive evidence that output elasticities vary systematically across firms, and show that they can lead to a spurious negative correlation between markups and firm size. I then show that Cobb-Douglas-implied markups decline with firm size, which is inconsistent with a body of evidence showing smaller price passthrough of cost shocks among larger firms (Amiti, Itskhoki, and Konings, 2014). I show that my estimates of translog output elasticities vary systematically across firms and are stable over time, leading to markup estimates that increase with firm size and are stable over time. Moreover, the estimated translog output elasticities closely track the input cost shares, which can be readily computed in the data and directly correspond to output elasticities under certain conditions.

I first examine whether output elasticities are likely to vary systematically across firms by looking at input cost shares. In my model in Section 2, the material and labor cost shares can be written as:

$$CS_{m,jt} = \frac{\alpha_{m,jt}}{\alpha_{k,jt} + \alpha_{h,jt} + \Lambda_{jt}\alpha_{h,jt}}$$
 and $CS_{h,jt} = \frac{\Lambda_{jt}\alpha_{h,jt}}{\alpha_{k,jt} + \alpha_{h,jt} + \Lambda_{jt}\alpha_{h,jt}}$

where the cost shares are measured as the expenditure on an input divided by total costs. Total costs are measured as the sum of material, labor, and capital expenditure, where capital expenditure is measured as expenditure on capital investments. These cost shares map directly into firm-specific output elasticities ($CS_{m,jt} = \alpha_{m,jt}$, $CS_{h,jt} = \alpha_{h,jt}$) if the following assumptions hold: constant returns to scale and perfectly competitive input markets. In my model, labor markets are imperfectly competitive, implying that observed input cost shares do not directly correspond to output

elasticities—disentangling output elasticities from labor wedges requires estimating production functions. Nevertheless, output elasticities are key components of input cost shares. Therefore, the evidence on cost shares presented in this section is suggestive of how actual output elasticities vary across firms.



Figure 7: Material and labor cost shares across firms.

This figure shows how material and labor cost shares vary across quantiles of firm size (sales or employment) and labor productivity, conditional on 2-digit sector×year fixed effects. Confidence intervals at the 95% level are plotted.

Figure 7 plots the material and labor cost shares against firm size (measured as sales and employment) and labor productivity (measured as sales per hour). Material cost shares increase systematically with firm size and labor productivity, while labor cost shares decrease with firm size and labor productivity. Moreover, Table 12 shows that there is substantial dispersion in cost shares within 2-digit sectors. These patterns suggest that material and labor output elasticities vary systematically across firms.

Cost shares	Mean	Median	10 th Pct	25 th Pct	75 th Pct	90 th Pct	Variance
Material	0.56	0.57	0.32	0.45	0.69	0.78	0.04
Labor	0.35	0.33	0.16	0.24	0.44	0.56	0.03

Table 12: The distribution of material and labor cost shares.

This table reports the summary statistics for the material and labor cost shares, conditional on 2-digit sector \times year fixed effects.

To the extent that the cross-sectional correlations between cost shares and firm size are driven by output elasticities, assuming Cobb-Douglas production functions will lead to systematic biases in the measurement of markups and labor wedges. To illustrate this, I write the relationship between Cobb-Douglas markups μ_{jt}^{CD} and "true" markups μ_{jt} (markups obtained without assuming homogenous output elasticities) as:

$$\mu_{jt}^{CD} = \mu_{jt} \frac{\alpha_m^{CD}}{\alpha_{m,it}}$$

where $\mu_{jt}^{CD} \equiv \alpha_m^{CD} \frac{P_{jt}Y_{jt}}{P_{m,t}M_{jt}}$ and $\mu_{jt} = \alpha_{m,jt} \frac{P_{jt}Y_{jt}}{P_{m,t}M_{jt}}$. The covariance between Cobb-Douglas markups and firm size can then be written as:

$$CV(\log \mu_{jt}^{CD}, \log size_{jt}) = CV(\log \mu_{jt}, \log size_{jt}) - CV(\log \alpha_{m,jt}, \log size_{jt})$$
(11)

Therefore, the covariance between Cobb-Douglas markups and firm size will be *downward-biased* if material output elasticities ($\alpha_{m,jt}$) increase with firm size.

Similarly, the relationship between Cobb-Douglas labor wedges and "true" labor wedges is:

$$\Lambda_{jt}^{CD} = \Lambda_{jt} \left(\frac{\alpha_m^{CD}}{\alpha_{m,jt}} \right) \left(\frac{\alpha_{h,jt}}{\alpha_h^{CD}} \right)$$

where $\Lambda_{jt}^{CD} \equiv \frac{\alpha_m^{CD}}{\alpha_h^{CD}} \frac{\Phi_{jt}H_{jt}}{P_{m,t}M_{jt}}$ and $\Lambda_{jt} = \frac{\alpha_{m,jt}}{\alpha_{h,jt}} \frac{\Phi_{jt}H_{jt}}{P_{m,t}M_{jt}}$. The covariance between Cobb-Douglas labor wedges and firm size can be written as:

$$CV(\log \Lambda_{jt}^{CD}, \log size_{jt}) = CV(\log \Lambda_{jt}, \log size_{jt}) - CV(\log \alpha_{m,jt}, \log size_{jt}) + CV(\log \alpha_{h,jt}, \log size_{jt})$$
(12)

Therefore, covariance between Cobb-Douglas labor wedges and firm size will also be *downward-biased* if material output elasticities increase with firm size, or if labor output elasticities decrease with firm size.

Figure 8 shows how different markup measures vary by firm size and labor productivity. The first column shows that Cobb-Douglas markups decline steeply with firm size and labor productivity, which is inconsistent with existing evidence on lower passthrough of cost shocks to prices among large firms (see, e.g., Amiti et al. (2014)). As shown above, this negative correlation may be due to Cobb-Douglas markups capturing a positive correlation between material output elasticities and firm size.

Without estimating production functions, I now attempt to assess the extent to

which Cobb-Douglas markup measures are likely to understate the relationship between actual markups and firm size. Suppose that material output elasticities are positively correlated with firm size. Then, equation (11) shows that controlling material output elasticities will increase the correlation between Cobb-Douglas markups and firm size. Although material output elasticities are not observed, under the assumption that production functions exhibit constant returns to scale, they depend on the capital-material and labor-material ratios, $\alpha_{m,jt} = \alpha_m (\frac{K_{jt}}{M_{jt}}, \frac{H_{jt}}{M_{jt}})$.²⁸ This suggests that the correlation between Cobb-Douglas markups and firm size would increase if one controls for these input ratios.



Figure 8: Cobb-Douglas and Translog (log) markup measures across firms.

This figure shows how different log markup measures vary by firm size (sales or employment) and labor productivity quantiles, conditional on 2-digit sector×year fixed effects. Confidence intervals at the 95% level are plotted.

The second column of Figure 8 shows that the slope between Cobb-Douglas markups and firm size becomes positive upon controlling for capital-material and labor-material ratios. The third column shows that markups implied by the estimated translog production functions are also increasing with firm size (without controlling for input ra-

²⁸Similarly, the labor output elasticity can be written as $\alpha_{h,jt} = \alpha_h (\frac{K_{jt}}{H_{jt}}, \frac{M_{jt}}{H_{jt}})$. Suppose labor output elasticities decrease with firm size, then equation (12) suggests that the correlation between Cobb-Douglas labor wedges and firm size will increase if one controls for the relevant input ratios.
tios). Figure 9 shows that a similar pattern holds for labor wedges: the Cobb-Douglas labor wedges decline with firm size more steeply than the translog counterpart.



Figure 9: Cobb-Douglas and Translog (log) labor wedge measures across firms.

This figure shows how different log labor wedge measures vary by firm size (sales or employment) and labor productivity quantiles, conditional on 2-digit sector \times year fixed effects. Confidence intervals at the 95% level are plotted.

Finally, I show that the estimated translog output elasticities vary across firms in a similar way compared to the input cost shares, and that they are stable over time. Table 13 shows that the translog material output elasticities are on average lower than their cost shares, while labor output elasticities are higher than their cost shares. This is what one would expect if labor wedges are below 1. Consistent with the fact that material cost shares increase with firm size depicted in Figure 7, Figure 10 shows that translog material output elasticities are also increasing with firm size, while these patterns are inverted for translog labor output elasticities. Further, Figure 11 shows that both the distribution of translog material and labor output elasticities are stable over time. This implies that the translog markup and labor wedge measures are also stable over time (see Figure 14).

Table 13: The distribution of output elasticities.

Elasticities	Mean	Median	10^{th} Pct	25 th Pct	75 th Pct	90 th Pct	Variance
Material	0.47	0.48	0.30	0.39	0.57	0.66	0.02
Labor	0.47	0.46	0.28	0.37	0.56	0.65	0.02
Capital	0.05	0.05	0.02	0.04	0.07	0.09	0.00

This table reports the summary statistics for the output elasticities, conditional on 2-digit sector \times year fixed effects.



Figure 10: Translog output elasticities across firms.

This figure shows how output elasticities vary across quantiles of firm size (sales or employment) and labor productivity, conditional on 2-digit sector×year fixed effects. Confidence intervals at the 95% level are plotted.



Figure 11: Translog output elasticities over time.

B.6.2 Leontief production function

Section 3.3 proposes a method for separately estimating worker bargaining power (κ) and firm monopsony power (λ_j). Identification of κ relies on the property that markups (μ_j) transmit into labor wedges (Λ_j) when workers have some bargaining power. Implementing this approach therefore requires firm-level estimates of markups and labor wedges. However, this approach does not hinge on a specific method for obtaining these estimates.

In Section 3.2, I obtain these estimates using the production-based approach following De Loecker and Warzynski (2012) and Yeh et al. (2022), assuming a *gross output* production function. Under this approach, firm markups and labor wedges are derived from the first-order conditions for materials and labor, given estimated output elasticities. These estimates are valid under relatively mild assumptions on preferences and product market competition.

A natural question is how my bargaining power estimation approach adapts when the production function is *Leontief* in materials—that is, when material inputs are nonsubstitutable. In such settings, the materials first-order condition does not hold, so the standard production approach is no longer applicable without further modifications. Alternative identification strategies are therefore needed to estimate markups and labor wedges.

Rubens (2023) offers important guidance in such cases. They show that, when production functions are Leontief in materials, the production approach can still deliver estimates of markups and markdowns if it is combined with either a specific model of (O1) input supply or (O2) output demand. In this section, I explain that, in my model, O1 is insufficient to separately identify labor wedges and markups. O2 is a more promising alternative for doing so, but it requires a different type of data and entails conceptual challenges that go beyond the scope of this paper.

Imposing a model of labor supply (O1). This option is not sufficient to separately identify markups and labor wedges—thus workers' bargaining power—because labor wedges differ from monopsony markdowns when $\kappa > 0$. To see this, consider the implications of a Leontief production function in materials. In this case, the labor

share of value-added implied by the model in Section 2 is:

$$\frac{\Phi_j H_j}{P_j Y_j} = \alpha_{h,j} \underbrace{\left(\kappa \left(1 - \frac{\alpha_{k,j}}{\mu_j}\right) \frac{\mu_j}{\alpha_{h,j}} + (1 - \kappa)\lambda_j\right)}_{=\Lambda_j} \mu_j^{-1}$$

where $P_j Y_j$ now denotes nominal value-added instead of revenues, and $(\alpha_{h,j}, \alpha_{k,j})$ are the labor and capital output elasticities of a value-added production function. Imposing a model of labor supply makes it possible to estimate labor supply elasticities and the implied monopsony markdowns (λ_j) . However, even with estimates of monopsony markdowns and output elasticities, the expression above shows that μ_j and Λ_j are still not separately identified without additional information on κ . In other words, imposing a labor supply model alone cannot disentangle the effects of product market power from those of worker bargaining power.

Imposing a model of output demand (O2). In contrast to O1, the labor wedge expression implies that, when production functions are Leontief in materials, labor wedges remain identified if markups and labor output elasticities are observed. This makes it possible to implement the estimation strategy for κ described in Section 3.3, even under non-substitutability in material inputs.

This alternative of strategy is therefore a promising path forward. Obtaining markup estimates by estimating demand systems has a long-standing tradition (Berry, Levinsohn, and Pakes, 1995). This demand-based approach provides estimates the own-price and cross-price demand elasticities, allowing researchers to compute markups.

However, estimating demand systems requires detailed product-level prices, quantities, and product characteristics, such as supermarket scanner data. Moreover, applying this demand-based approach to estimate markups for manufacturing firms requires modeling the mapping between the prices charged by producers and the prices paid by final consumers. I refer interested readers to De Loecker and Scott (2024) for a detailed discussion on the challenges involved when combining the production-based and demand-based approaches to estimate markups.

B.7 The choice of flexible inputs

The production function approach to markup estimation requires selecting a flexible input to measure markups with. When implementing this approach, I assume that materials are flexible inputs, as is standard in the literature (De Loecker and Warzynski, 2012). However, another natural candidate for the flexible input is electricity consumption. Relatedly, electricity prices may be less subject to potential monopsony power by buyer firms, compared to materials. However, after analyzing firm-level electricity consumption data from the French Survey of Energy Consumption in Industry (EACEI), I find that the measurement error in electricity consumption is likely much larger than in materials. Combined with the small share of electricity expenditures in revenue, this leads to highly dispersed and unreliable markup estimates.

The EACEI survey data provides information on energy consumption and purchase for French manufacturing firms from 1983 to 2021. The survey includes approximately 5,000 firms per year. For each year and plant combination, I observe the total amount of electricity consumed (in megawatt hours, MwH) and total expenditure on electricity (in Euros). I aggregate the data to the firm-year level. I drop firms with missing or negative values of electricity consumption and expenditure. I then link this survey to my estimation sample (EAP-FARE 2009-2016). However, this significantly reduces the sample size in my estimation sample to 28,081.

		Materia	s	Electricity			
	Mean	Median	Variance	Mean	Median	Variance	
Cost share	0.58	0.60	0.03	0.02	0.01	0.00	
Revenue share	0.40	0.41	0.03	0.01	0.01	0.00	
log <u>Revenue</u>	1.06	0.90	0.45	4.90	4.90	0.99	

Table 14: The cost and revenue shares of electricity and materials.

This table reports the cost and revenue shares of electricity. In this table, costs are measured as the sum of wage bills, material expenditure, capital investments, and electricity expenditure. The variances are across firms within 2-digit sectors. The sample consists of firms that are surveyed in the EACEI energy consumption data and the EAP-FARE output price and balance sheet data.

The survey reveals that electricity expenditure is a small share of total costs and revenue among French manufacturing firms. Table 14 shows that the cost share and revenue share of electricity is about 1% at the median French firm in this sample. In comparison, the cost and revenue shares of material inputs are many times larger, at

60% and 41%.

This small revenue share of electricity presents a challenge for using it as a flexible input to measure markups. Even a small amount of measurement error in electricity expenditure can cause large variability in the resulting markup estimates, since markups are calculated as the output elasticity of the flexible input multiplied by the inverse of its revenue share.

Table 14 shows that the revenue-to-electricity-expenditure ratio (in logs) is much more dispersed than the revenue-to-material expenditure ratio. Moreover, the last row of Table 15 shows that electricity consumption growth is much more volatile than materials growth. When examining the comovement of electricity consumption with other factor inputs, I find that materials exhibit a stronger comovement with capital and labor compared to electricity consumption (see the third and fourth rows of Table 15). Taken together, these descriptive evidence is suggestive of larger measurement error in electricity expenditure.

	%∆ Capital	%Δ Labor	$\%\Delta$ Materials	%Δ Electricity
%Δ Capital	1.00			
%∆ Labor	0.13	1.00		
$\%\Delta$ Materials	0.11	0.09	1.00	
$\%\Delta$ Electricity	0.04	0.04	0.06	1.00
Variance	0.02	0.03	0.10	0.18

Table 15: Firm-level comovement of factor inputs.

This table reports the Pearson correlation coefficients between the firm-level growth rates of capital, effective labor, materials, and electricity consumption.

B.8 Monopsony in the market for materials

In Section 3, I measure markups and labor wedges under the assumption that firms have no monopsony power in material inputs. If firms have monopsony power in material markets, then even if one has data on actual material and labor output elasticities, the material and labor first-order conditions do not deliver the correct measures of markups and labor wedges; they will be contaminated by a material monopsony wedge. These first-order conditions become:

$$\mu_{jt} = \lambda_{m,jt} \alpha_{m,jt} \frac{P_{jt} Y_{jt}}{P_{m,jt} M_{jt}} \quad \text{and} \quad \Lambda_{jt} = \lambda_{m,jt} \frac{\alpha_{m,jt}}{\alpha_{h,jt}} \frac{\Phi_{jt} H_{jt}}{P_{m,jt} M_{jt}}$$

where $\lambda_{m,jt} = \frac{\epsilon_{m,jt}}{1+\epsilon_{m,jt}}$ is the monopsony markdown for materials, which depends on the material supply elasticity $\epsilon_{m,jt}$.

I now show that, in the presence of material monopsony power, it is still possible to obtain correct measures of markups and labor wedges—not contaminated by material markdowns—if production functions are estimated using *material expenditure* instead of quantities. Without material price data, one recovers the output elasticity with respect to material expenditure when estimating production functions. In what follows, I derive this material expenditure output elasticity and show that (a) it is equal to the composite elasticity $\tilde{\alpha}_{m,jt} \equiv \lambda_{m,jt} \alpha_{m,jt}$ when firms have material monopsony power, and (b) it is equal to the material quantity output elasticity $\alpha_{m,jt}$ when firms are pricetakers in material markets. Therefore, when firms have material monopsony power, the correct markup and labor wedge measures can still be obtained, although one cannot disentangle material monopsony power from material (quantity) output elasticities.

Consider the following upward-sloping material supply curve: $M_{jt} = \mathcal{V}(P_{m,jt})$. The inverse material supply curve, expressed as a function of material expenditure, is: $P_{m,jt} = \tilde{\mathcal{V}}(\tilde{M}_{m,jt})$, where $\tilde{M}_{jt} \equiv P_{m,jt}M_{jt}$. Variables written in lowercase letters are in logs. The production function takes the form:

$$y_{jt} = f(k_{jt}, h_{jt}, m_{jt}) + \omega_{jt}$$
$$= f(k_{jt}, h_{jt}, \tilde{m}_{jt} - \tilde{v}(\tilde{m}_{jt})) + \omega_{jt}$$

where the second equality uses the inverse material supply curve. The output elasticity with respect to material expenditure is:

$$\tilde{\alpha}_{m,jt} \equiv \frac{\partial y_{jt}}{\partial \tilde{m}_{jt}} = \alpha_{m,jt} \left(1 - \frac{\partial \tilde{v}(\tilde{m}_{jt})}{\partial \tilde{m}_{jt}} \right)$$

What remains to be shown is that the term $\left(1 - \frac{\partial \tilde{v}(\tilde{m}_{jt})}{\partial \tilde{m}_{jt}}\right)$ is equal to the material markdown. From the inverse material supply curve, the elasticity of material prices with respect to material expenditure is:

$$\frac{\partial p_{m,jt}}{\partial \tilde{m}_{jt}} = \frac{\partial p_{m,jt}}{\partial m_{jt}} \left(\frac{\partial \tilde{M}_{jt}}{\partial M_{jt}}\right)^{-1} P_{m,jt} \tag{B6(a)}$$

and the partial derivative of material expenditure with respect to material quantities

is:

$$\frac{\partial \tilde{M}_{jt}}{\partial M_{jt}} = \left(1 + \left(\underbrace{\frac{\partial m_{jt}}{\partial p_{m,jt}}}_{\equiv \epsilon_{m,jt}}\right)^{-1}\right) P_{m,jt} \tag{B6(b)}$$

Combining equations (B6(a)) and (B6(b)) gives $\lambda_{m,jt} \equiv \frac{\epsilon_{m,jt}}{1+\epsilon_{m,jt}} = \left(1 - \frac{\partial \tilde{v}(\tilde{m}_{jt})}{\partial \tilde{m}_{jt}}\right)$. Therefore, the output elasticity with respect to material expenditure is equal to the composite elasticity $\tilde{\alpha}_{m,jt} \equiv \lambda_{m,jt} \alpha_{m,jt}$. When firms have no material monopsony power, so that $\lambda_{m,jt} = 1$ and $P_{m,jt} = P_{m,t}$, the output elasticity with respect to material expenditure or quantities are the same. This result shows that estimating production functions based on material expenditures rather than quantities recovers the output elasticity with respect to material expenditure, which hinders the measurement of material markdowns, but not markups or labor wedges.

The presence of material markdowns has implications for the use of material inputs as a proxy variable when estimating production functions. Given the material supply equation faced by the firm, the material input FOC can be written as:

$$\tilde{M}_{jt} = \left(1 - \frac{\partial \tilde{v}(\tilde{m}_{jt})}{\partial \tilde{m}_{jt}}\right) P_{jt} \mu_{jt}^{-1} \alpha_{m,jt} \Omega_{jt} F\left(K_{jt}, H_{jt}, \frac{\tilde{M}_{jt}}{\tilde{\mathcal{V}}(\tilde{M}_{m,jt})}\right)$$

Thus, the control function for productivity (Ω_{jt}) is formed on material expenditure rather than quantities. The presence of material markdowns do not violate the scalar unobservable assumption if firm size (expenditure on materials) is a sufficient statistic for material markdowns, analogous to the implication that firm size in terms of wagebill shares are sufficient statistics for labor market monopsony markdowns in recent models such as Berger et al. (2022).

In deriving these results, I abstract from material price differences driven by quality heterogeneity. In the presence of quality heterogeneity, there will be material price dispersion across firms even when firms do not have material monopsony power. The lack of material price data will cause an input price bias, as discussed in Section 3. De Loecker et al. (2016) develop an approach to addressing input price bias using output price data, and I implement their approach to address this concern.

B.9 Alternative methods for production function estimation

The production approach to markup measurement requires estimating the output elasticity of a flexible input (De Loecker and Warzynski, 2012) (DLW henceforth). In Section 3, I estimate a translog gross output production function to obtain the flexible input's (materials) output elasticity. As Bond and Söderbom (2005) and Gandhi et al. (2020) show, this output elasticity is challenging to identify, particularly in the context of gross output production functions and perfectly competitive goods markets. In such cases, persistence in material prices is the only source of persistence in material input choices, which is crucial for the identification of the material output elasticity. However, De Ridder et al. (2021) show that with imperfectly competitive product markets—a key feature of my model—variation in output prices and productivity shocks to firms' competitors generate additional variation in firms' material input demand, which aids the estimation of the material output elasticity.

To more directly address the identification concern raised in Gandhi et al. (2020), I compare my baseline estimates of output elasticities to those estimated using two alternative methods that aim to address this identification challenge: (i) control function method with constant returns to scale production functions (Flynn, Gandhi, and Traina, 2019), and (ii) dynamic panel method (Blundell and Bond, 1998).

	DLW			FGT			BB		
	Mean	Median	Var	Mean	Median	Var	Mean	Median	Var
Material	0.47	0.48	0.03	0.48	0.49	0.04	0.36	0.37	0.02
Labor	0.47	0.46	0.03	0.43	0.42	0.04	0.34	0.33	0.02
Capital	0.05	0.05	0.00	0.08	0.08	0.00	0.07	0.07	0.00
Returns to scale	0.99	0.98	0.00	1.00	1.00	0.00	0.78	0.79	0.02

Table 16: Output elasticities under alternative production function estimation methods.

This table reports the estimated output elasticities under alternative production function approaches. My baseline approach follows De Loecker and Warzynski (2012), referred to as DLW. This table also presents estimates from implementing constant returns to scale (Flynn et al. (2019); FGT) and dynamic panel methods (Blundell and Bond (1998); BB). All output elasticities are trimmed by 1% on either side of their distribution.

In the context of imperfectly competitive output markets, Flynn et al. (2019) (FGT henceforth) show that the challenge of identifying the material output elasticity can be recast as a challenge of separately identifying returns to scale from the level of markups. They further show that under the restriction of constant returns to scale, the

material output elasticity is identified. Table 16 compares my baseline estimates of the output elasticities to those obtained under constant returns to scale. Comparing the first three columns with the middle three columns shows that my baseline estimates (based on DLW) and those using FGT are similar. Figure 12 shows that material output elasticities estimated with either DLW or FGT both increase with firm size (as defined by sales). The figure also shows that the implied markups and labor wedges increase with firm size. Figures 13 and 14 further show that the estimated material output elasticities, markups, and labor wedges are stable over time.



Figure 12: Output elasticities, markups, and labor wedges by firm size under DLW and FGT.

This figure compares the estimated output elasticities, markups, and labor wedges using two methods: (i) my baseline approach following De Loecker and Warzynski (2012), and (ii) the Flynn et al. (2019) approach which imposes a constant returns to scale restriction. Firm size quantiles are defined based on sales. The plots control for 2-digit sector×year fixed effects. Vertical bars are 95% confidence intervals.

As an alternative to the control function method, the dynamic panel data method of Blundell and Bond (1998) (BB henceforth) circumvents the need to select a proxy variable, instead relying on a log-linear AR(1) process for Hicks-neutral productivity to identify production function parameters. In practice, the estimated output elasticities and returns to scale are much lower compared to my baseline estimates, as shown in Table 16. This leads to markup measures that fall below 1 for most firms in my sample. Because the BB method requires first-differencing input and output data (and their lags), it significantly reduces variation in the data that can be used for identification since cross-sectional variation in factor inputs are differenced away. Further, attenuation bias due to measurement error is exacerbated in dynamic panel methods. In a Monte Carlo simulation comparing control function methods and dynamic panel methods when the true production function has constant returns to scale, Yeh et al. (2022) also find that dynamic panel methods tend to estimate far lower returns to scale. The practical difficulty of obtaining reasonable estimates of production function function parameters using dynamic panel methods is also discussed in De Loecker and Syverson (2021).



Figure 13: Time-series of material output elasticities under DLW and FGT.



Figure 14: Time-series of markups and labor wedges under DLW and FGT.

B.10 Multiplication bias

The key estimating equation (3) for worker bargaining power requires first measuring labor wedges (Λ) and product market rents ($\tilde{\mu}$). Measuring them requires data on material expenditure as well as estimates of material and labor output elasticities (α_m, α_h). Because these variables appear on both sides of the estimating equation multiplicatively, any measurement error or estimation bias in these variables will be positively correlated and, therefore, lead to a mechanical upward bias of the bargaining power estimate $\hat{\kappa}$. I refer to this bias as *multiplication bias*.

There are two reasons why the output elasticities could be measured with error, conditional on the production function specification. First, material and labor output elasticities are functions of factor inputs, and these inputs could be measured with error. Second, the parameters governing the material and labor output elasticities could be estimated with a bias due to the difficulties of identifying the flexible input's output elasticity (Gandhi et al., 2020).

To see how biased or mismeasured output elasticities can affect $\hat{\kappa}$, consider the following example. As a reminder, I denote measured or estimated parameters or variables with a hat. Let the true markdown be constant and homogenous, $\lambda_{jt} = \lambda$. Let ϵ_{jt}^{bias} be the source of bias (such as measurement error or estimation biases), such that $\frac{\hat{\alpha}_{m,jt}}{\hat{\alpha}_{h,jt}} = \frac{\alpha_{m,jt}}{\alpha_{h,jt}} \epsilon_{jt}^{bias}$. Let $\tilde{\epsilon}_{jt}$ be a random error uncorrelated with $\hat{\mu}_{jt}$. The estimating equation becomes:

$$\underbrace{\frac{\Phi_{jt}H_{jt}}{P_{m,t}M_{jt}}\frac{\hat{\alpha}_{jt}^{m}}{\hat{\alpha}_{jt}^{h}}}_{\hat{\Lambda}_{jt}} = \kappa \underbrace{\left(\frac{P_{jt}Y_{jt} - P_{m,t}M_{jt} - P_{k,t}K_{jt}}{P_{m,t}M_{jt}}\right)\frac{\hat{\alpha}_{jt}^{m}}{\hat{\alpha}_{jt}^{h}}}_{\hat{\mu}_{jt}} + (1-\kappa)\lambda + \tilde{\epsilon}_{jt}$$

Then, an OLS regression of $\hat{\Lambda}_{jt}$ on $\hat{\mu}_{jt}$ will give an upward-biased estimate of κ :

$$\hat{\kappa}^{OLS} = \kappa + (1 - \kappa) \frac{CV(\lambda \epsilon_{jt}^{bias}, \tilde{\mu}_{jt} \epsilon_{jt}^{bias})}{V(\tilde{\mu}_{jt} \epsilon_{it}^{bias})} \ge \kappa, \text{ where } \kappa \in [0, 1]$$

Therefore, the upward bias is larger when the true κ is smaller, and one may estimate positive worker bargaining power even when $\kappa = 0.29$

²⁹Here, I have assumed that $\lambda_j = \lambda$, so there is only one source of bias—the multiplication bias, which biases \hat{k} upwards. If λ_j is endogenous and depends on firm size and amenities, then monopsony

In the rest of this section, I discuss several attempts to address this concern, depending on the source of the bias (ϵ^{bias}). First, under a classical measurement error assumption, I obtain an upper bound estimate of the measurement error in $(\hat{\alpha}_m / \hat{\alpha}_h)$ and simulate the extent of bias on $\hat{\kappa}$. Second, I follow a well-established tradition of using instrumental variables to address measurement error. In the same simulation, I show that instrumenting product market rents using their lags $(\hat{\mu}_{it-1})$ recovers the true κ . I therefore also report estimates when product market rents are instrumented using their lags. However, this lags-as-IV strategy does not address the multiplication bias when ϵ^{bias} are serially correlated. This is the case when the parameters governing the output elasticities are estimated with bias. This is a concern since the production function estimation literature has shown that identifying the output elasticities of flexible inputs is not trivial (see, for example, Bond and Söderbom (2005) and Gandhi et al. (2020)). To address this concern, I propose an alternative estimating equation that avoids having to first estimate output elasticities. However, as I explain later in this section, this alternative requires approximating a more complex, potentially highly nonlinear function.

Measurement error and instrumental variables. Assuming that the source of bias on $\hat{\kappa}$ is classical measurement error in the output elasticities, I now use the panel structure of the data to infer the size of the measurement error. I then use this variance to simulate the likely extent of bias for a given distribution of $\tilde{\mu}_{it}$, and known λ and κ .

Let the log of the ratio of material-to-labor output elasticities be $x_{jt} \equiv \log \frac{\alpha_{m,jt}}{\alpha_{h,jt}}$, and let it be measured with error: $\hat{x}_{jt} = x_{jt} + \varepsilon_{jt}^{me}$, where \hat{x}_{jt} is the observed measure and $\varepsilon_{jt}^{me} = \log \varepsilon_{jt}^{me}$ is classical measurement error. Following Krueger and Summers (1988), I compute the within-firm variance as $V_j(\hat{x}_{jt}) = V_j(x_{jt}) + V(\varepsilon_{jt}^{me})$. Suppose that x_{jt} is highly persistent (or fixed) over time, so that $V_j(x_{jt}) \approx 0$. Then, taking the average of the firm-level variance gives $E[V_j(\hat{x}_{jt})] = V(\varepsilon_{jt}^{me})$. This provides an upper bound estimate of the extent of measurement error, because any within-firm variation in x_{jt} will be attributed to $V(\varepsilon_{jt}^{me})$.

Given a measure of the size of the measurement error, I assess the potential size of the multiplication bias. I simulate the relationship between labor wedges and product

markdowns bias $\hat{\kappa}$ downwards when firms with high product market rents (higher $\tilde{\mu}_j$) also have more monopsony power (lower λ_j). While the two sources of bias work in opposite directions, there is no reason *a priori* to expect them to cancel out.

market rents for a known κ and estimate κ with and without the presence of measurement error. For this purpose, I continue to assume that monopsony markdowns are constant and homogenous across firms, and calibrate it to $\lambda = 0.5$. I then take the estimated distribution product market rents as the true distribution and compute the implied Λ_{jt} under different values of κ . Next, I draw a sequence of ε_{jt}^{me} from $\mathcal{N}(0, V(\varepsilon_{jt}^{me}))$, where $V(\varepsilon_{jt}^{me})$ is estimated from the previous step, and multiply $\tilde{\mu}_{jt}$ and Λ_{jt} by ε_{jt}^{me} . I then compare the difference between the true κ and those estimated by OLS (with and without the presence of measurement error). Without measurement error, the OLS estimator for κ is unbiased under the assumption of a constant and homogenous λ .

Table 17 compares the true κ and the OLS estimate of κ for different sizes of measurement error. The estimate of the size of measurement error following Krueger and Summers (1988) is $V(\varepsilon^{me}) = 0.102$. For measurement error of this size, panel (a) shows that the OLS estimate of κ is close to the true κ , with an upward bias of around 0.007. Panels (b) and (c) show that for much larger measurement error the bias is also much larger, although the size of the bias declining in the level of κ .

		True κ							
	0.05	0.10	0.15	0.20	0.25	0.30	0.35		
Panel (a)			$V(\epsilon)$	^{ne})=0.102					
OLS	0.058	0.107	0.157	0.206	0.256	0.306	0.355		
IV	0.050	0.100	0.150	0.200	0.250	0.300	0.350		
Panel (b)			$V(\epsilon$	^{.me})=0.20					
OLS	0.076	0.125	0.173	0.222	0.270	0.319	0.368		
IV	0.051	0.100	0.151	0.200	0.251	0.300	0.350		
Panel (c)		$V(\epsilon^{me})=0.30$							
OLS	0.101	0.148	0.195	0.243	0.290	0.337	0.384		
IV	0.050	0.100	0.150	0.200	0.250	0.300	0.350		

Table 17: Effect of measurement error on estimated worker bargaining power.

This table compares the true and estimated workers' bargaining power parameter in the presence of measurement error in material expenditures or output elasticities.

A well-established approach to addressing measurement error in the explanatory variable is to use instrumental variables (IV). The instrument needs to be correlated with the explanatory variable and be orthogonal to the measurement error. Under a classical measurement error assumption, product market rents can be instrumented using their lags. Table 17 compares the OLS and IV estimates of κ , where the the IV is

lagged product market rent. In all three panels, the instrument successfully recovers the true κ .

An alternative estimating equation for worker bargaining power. The strategy above of using lagged product market rent as an IV to address the multiplication bias induced by measurement error relies on the assumption of classical measurement error. This is a strong assumption; if the production function parameters are not identified, then the bias in the estimated output elasticities will be serially correlated. This causes $\hat{\kappa}$ to be upward-biased and lagged product market rent is no longer a valid instrument. To avoid having to first estimate the output elasticities (α), I propose using an alternative estimating equation to (3). I rearrange the baseline estimating equation (3) to obtain:

$$\frac{\Phi_{jt}H_{jt}}{P_{m,t}M_{jt}} = \kappa \left(\frac{P_{jt}Y_{jt} - P_{m,t}M_{jt} - P_{k,t}K_{jt}}{P_{m,t}M_{jt}}\right) + (1 - \kappa)\lambda(H_{jt}, \mathcal{H}^{-1}(H_{jt}, \Phi_{jt}H_{jt}))\frac{\alpha_h(K_{jt}, M_{jt}, H_{jt})}{\alpha_m(K_{jt}, M_{jt}, H_{jt})}$$
(13)

Compared to equation (3), estimating equation (13) does not require estimates of output elasticities. However, equation (13) includes a more complicated term to approximate $(\lambda(.)\alpha_h(.)/\alpha_m(.))$.³⁰ In practice, I approximate $\lambda(.)\alpha_h(.)/\alpha_m(.)$ with a fourth-order polynomial. Similar to the baseline equation (3), this alternative equation is also not immune to measurement-error-induced multiplication bias because material expenditure and wage bills appear on both sides of the equation (13), I also instrument all right-hand-side variables using their lags under a classical measurement error assumption.

Table 18 presents the bargaining power estimates from implementing equation (13). Columns (1) through (3) show that controlling for the term $\lambda(.)\frac{\alpha_h(.)}{\alpha_m(.)}$ and firm fixed effects are important, as they substantially reduce the estimated bargaining power. Indeed, column (4) shows that a specification that controls for both $\lambda(.)\frac{\alpha_h(.)}{\alpha_m(.)}$ and firm fixed effects gives an estimated κ of 0.149, close to my baseline estimate of 0.125 in column (4) of Table (2). Column (6) shows that additionally using lagged left-hand-side

³⁰Note that, if the production function were Cobb-Douglas, then we need only approximate $\lambda(.)$, which would make equation (13) more attractive than equation (3).

variables to address multiplication bias due to classical measurement error further reduces the estimated κ to 0.137.

	$rac{\Phi_{jt}H_{jt}}{P_{m,t}M_{it}}$							
	(1)	(2)	(3)	(4)	(5)	(6)		
$\frac{P_{jt}Y_{jt} - P_{m,t}M_{jt} - P_{k,t}K_{jt}}{P_{m,t}M_{it}}$	0.248	0.248	0.131	0.149	0.129	0.137		
····	(0.012)	(0.010)	(0.007)	(0.007)	(0.009)	(31.563)		
Sector×year fixed effects	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark		
Firm fixed effects		\checkmark		\checkmark		\checkmark		
Control for $\lambda(.) \frac{\alpha_h(.)}{\alpha_m(.)}$			\checkmark	\checkmark	\checkmark	\checkmark		
IV for measurement error					\checkmark	\checkmark		
Observations	104,865	101,941	104,865	101,941	81,482	78,818		

Table 18: Estimated worker bargaining power using an alternative equation.

This table reports the estimated worker bargaining power parameter using the alternative estimating equation (13). When controlling for the term $\lambda(.)\frac{\alpha_h(.)}{\alpha_m(.)}$, I use a second-order polynomial of its arguments: capital, materials, effective labor, and wage bills. Bootstrapped standard errors are reported in parentheses. In columns (5) and (6), all left-hand-side variables are instrumented with their lags.

Taken together, these findings suggest that measurement error in factor inputs and estimation biases due to challenges in identifying production function parameters do not result in significant upward bias in my estimates of worker bargaining power. Nevertheless, it is useful to compare how measured monopsony markdowns depend on the value of inferred worker bargaining power.

Measured monopsony markdowns under different values of worker bargaining power. To compare measured markdowns under different values of κ , I start by deriving an upper and lower bound for κ . Given the estimated $\tilde{\mu}_{jt}$ and Λ_{jt} , we can obtain upper and lower bounds for κ by noting that $\lambda_{jt} \in [0, min\{1, \Lambda_{jt}\}]$. The labor wedge equation (3) can be rearranged as follows:

$$\kappa = rac{\Lambda_{jt} - \lambda_{jt}}{ ilde{\mu}_{jt} - \lambda_{jt}}$$

In this expression, κ is decreasing in λ . Since the maximum value λ_{jt} can take is Λ_{jt} , the lower bound for κ is 0. Since the minimum value λ_{jt} can take is 0, the upper bound for κ is $\frac{\Lambda_{jt}}{\tilde{\mu}_{jt}}$. Because firms are assumed to have the same bargaining parameter, I set the upper bound as the level of κ such that $\lambda_{jt} = 0$ for at most a quarter of firms. The

bounds on worker bargaining power are then $\kappa \in [0, 0.28]$.

Table 19 shows the summary statistics of measured markdowns for different values of κ . The table shows that the mean and median markdowns decrease steeply with κ . Therefore, the larger the estimated κ is, the greater the labor market power firms are inferred to have.

		Measured λ								
κ	0.05	0.10	0.15	0.20	0.25					
Mean	0.54	0.47	0.40	0.31	0.21					
Median	0.55	0.48	0.41	0.33	0.25					
25 th percentile	0.46	0.39	0.30	0.20	0.09					
75 th percentile	0.64	0.58	0.52	0.45	0.38					
75-25 difference	0.18	0.19	0.22	0.25	0.31					

Table 19: Measured monopsony markdowns under different κ .

This table compares the distribution of measured monopsony markdowns under different values of worker bargaining power κ .



Figure 15: Implied monopsony markdowns by firm wage premia and firm size.

How does the value of κ affect the correlation between measured monopsony markdowns and firm characteristics? The first panel of Figure 15 shows that the correlation between markdowns and firm wage premia is relatively flat across different values of κ , although this correlation flips from positive to negative as κ increases. The second panel of Figure 15 shows that measured markdowns decline with firm size for each value of κ , although the relationship between monopsony markdowns and firm size is also relatively flat.

On the determinants of the implied monopsony markdown. One observation from Figure 15 is that the implied monopsony markdown widens sharply with the value of worker bargaining power. To clarify this pattern, I describe how each component of the labor wedge equation affects: (a) the level of the implied monopsony markdown, and (b) the sensitivity of the implied markdown with respect to bargaining power.

In the baseline model in Section 2, workers' collective outside options are assumed to be zero wages in the event of a (temporary) strike. To explore the implications of this assumption, consider a more general version of the collective bargaining problem in which the outside option is firm-specific and potentially non-zero. The problem becomes:

$$\max_{\Phi_j, P_j, M_j, K_j} \left((\Phi_j - \Phi_j^o) H_j \right)^{\kappa} \left(\Pi_j \right)^{1-\varepsilon}$$

subject to the labor supply curve $H_j = \mathcal{H}(\Phi_j - \Phi_j^o, A_j)$, product demand curve $Y_j = \mathcal{G}(P_j, D_j)$, and production technology $Y_j = \Omega_j F(K_j, M_j, H_j)$, with firm profits $\Pi_j = P_j Y_j - \Phi_j H_j - P_m M_j - P_k K_j$. Let Φ_j^o denote an arbitrary, non-zero, firm-specific outside option faced by workers employed by firm *j*. In this setup, the labor wedge becomes:

$$\Lambda_{j} = \frac{\kappa \widetilde{\mu}_{j} + (1 - \kappa)\lambda_{j}}{1 - (1 - \kappa)(1 - \lambda_{j})\frac{\Phi_{j}^{o}}{\Phi_{j}}} \quad \text{or, rearranging:} \quad \lambda_{j} = \frac{\Lambda_{j} \left(1 - (1 - \kappa)\frac{\Phi_{j}^{o}}{\Phi_{j}}\right) - \kappa \widetilde{\mu}_{j}}{(1 - \kappa)\left(1 - \frac{\Phi_{j}^{o}}{\Phi_{j}}\right)}$$

As before, $\tilde{\mu}_j$ are product market rents, which depend on markups. When worker outside wages are zero ($\Phi_j^o = 0$), this labor wedge expression becomes equation (3).

This labor wedge expression shows how product market rents, outside options, and worker bargaining power affect the level of the implied λ_i , all else equal:

- A higher worker bargaining power (κ) implies a lower λ_i , i.e., wider markdown.
- ▶ A higher markup—hence, product market rent ($\tilde{\mu}_i$)—implies a wider markdown.
- A higher outside option $\left(\frac{\Phi_j^o}{\Phi_j}\right)$ also implies a wider markdown.

To understand the sharp widening in implied monopsony markdowns when κ increases, as shown in Figure 15, it is useful to write the partial derivative of λ with respect to κ :

$$\frac{\partial \lambda_j}{\partial \kappa} = -\left(1 - \frac{\Phi_j^o}{\Phi_j}\right) (\widetilde{\mu}_j - \Lambda_j) \le 0$$

This derivative is negative since $\Lambda_j \leq \tilde{\mu}_j$ and $\Phi_j^o \leq \Phi_j$. This expression shows that:

- A higher markup—hence, μ̃_j—increases the impact of κ on the implied λ_j.
 A higher outside option (Φ^o_j/Φ_j) reduces the impact of κ on the implied λ_j.

The direction of the effect of workers' collective outside options on the level of the implied markdowns is therefore ambiguous. On one hand, higher outside options widen the implied markdowns (holding other parameters constant); on the other hand, they dampen the sensitivity of the implied markdowns to bargaining power.

This ambiguity raises the conceptual question: how should workers' outside options be modeled? What gives workers the leverage to demand wages above the monopsony level? One natural approach is to model outside offers as arriving through on-the-job search, allowing workers to use outside options to renegotiate wages when they arise (Cahuc et al., 2006). However, this leads to a model of *individual* wage bargaining in which outside options vary within firms. I discuss this alternative in Appendix C.1.

B.11 Labor adjustment costs

The estimated labor wedge $\hat{\Lambda}_{jt}$ could capture labor adjustment costs. I now derive labor wedges in the presence of labor adjustment costs and discuss how labor adjustment costs could bias the estimated bargaining power $\hat{\kappa}$ and the inferred monopsony markdown $\hat{\lambda}_{jt}$. I then attempt to assess the magnitude of this source of bias and incorporate labor adjustment costs in my estimation of κ . In this section, I will use the term 'employment' to mean effective labor (H_{jt}).

Specification of labor adjustment costs. There are two alternatives—convex and nonconvex cost functions—and they have different implications for the observed distribution of employment growth. The former implies that employment growth rates should be smooth and centered around zero. The latter implies that employment growth rates should be multimodal—most firms do not adjust employment, but when they do, they tend to make lumpy adjustments. In the French data, employment growth rates are smoothly distributed around zero (see Figure 16), hence I specify convex labor adjustment costs. Specifically, I work with quadratic adjustment costs in the derivations of labor wedges below.



Figure 16: Distribution of employment growth rates.

A deterministic model of firm growth with labor adjustment costs. I now use a stylized model to show that labor adjustment costs: (i) imply that labor wedges depend on firms' current and future employment growth rates; (ii) imply that one will find $\Lambda_j \in [0, 1]$ even if product and labor markets were perfectly competitive; (iii) disappear as firms grow towards their optimal sizes—that is, Λ_j grows as firms grow. For simplicity, suppose that firms face a constant elasticity goods demand function $P_{jt} = Y_{jt}^{-1/\sigma}$, a constant elasticity labor supply curve $\Phi_{jt} = H_{jt}^{1/\eta}$, and operate a production technology that uses only labor inputs $Y_{jt} = \Omega_j H_{jt}$. Further, assume that firm productivity Ω_j is heterogeneous across firms but constant over time. Consider the following wage bargaining problem:

$$\max_{H_{jt}} \left(\Phi_{jt} H_{jt} \right)^{\kappa} \left(V(\Omega_j, H_{jt-1}) - V^o(\Omega_j, H_{jt-1}) \right)^{1-\kappa}$$
(14)

subject to the labor supply curve and production technology. The firm's surplus is $V(\Omega_j, H_{jt-1}) - V^o(\Omega_j, H_{jt-1})$. If workers and firms reach an agreement on the terms of employment, the value function of the firm is:

$$V(\Omega_j, H_{jt-1}) = P_{jt}Y_{jt} - \Phi_{jt}H_{jt} - C(H_{jt}, H_{jt-1}, \Phi_{jt}) + \beta V(\Omega_j, H_{jt})$$

where β is the discount factor and the quadratic adjustment costs are $C(H_{jt}, H_{jt-1}, \Phi_{jt}) = \frac{\gamma}{2} \left(\frac{H_{jt}}{H_{jt-1}} - 1\right)^2 \Phi_{jt} H_{jt}$. If workers and firms do not arrive at an agreement, the firm's value function becomes:

$$V^{o}(\Omega_{j}, H_{jt-1}) = \beta V(\Omega_{j}, H_{jt}^{o})$$

where H_{jt}^{o} are the workers that firms carry over from time t to t + 1 if no agreement was reached at t. I assume that $H_{jt}^{o} = H_{jt-1}$, the stock of workers available to the firm at t + 1 if workers and firms do not agree at t is equal to employment at t - 1, so no new hires are added at t. Workers supply no labor at t if an agreement is not reached at t, and so there is no production, no wage expenditure, and no adjustment cost in that time. The FOC with respect to labor combined with the envelope theorem gives:

$$\Phi_{jt} = \frac{\mu^{-1}\Omega_{j}^{\frac{\sigma-1}{\sigma}}H_{jt}^{-\frac{1}{\sigma}} + \beta \frac{V_{jt} - V_{jt}^{o}}{H_{jt}}}{\left[\left(1 + \frac{\gamma}{2}g_{h,jt}^{2}\right) + (1 - \kappa)\lambda\left(\gamma g_{h,jt}(1 + g_{h,jt}) - \beta \gamma g_{h,jt+1}(1 + g_{h,jt+1})^{2}(1 + g_{\Phi,jt+1})\right)\right]}$$

where the firm's surplus can be written as:

$$V_{jt} - V_{jt}^{o} = P_{jt}Y_{jt} - \Phi_{jt}H_{jt} - C(H_{jt}, H_{jt-1}, \Phi_{jt}) + \beta \left[V(\Omega_{j}, H_{jt}) - V(\Omega_{j}, H_{jt}^{o})\right]$$

To simplify the FOC, I approximate the last term in the firm's surplus as follows:

$$V(\Omega_{j}, H_{jt}) - V(\Omega_{j}, H_{jt}^{o}) \approx \frac{\partial V(\Omega_{j}, H_{jt}^{o})}{\partial H_{jt}^{o}} (H_{jt} - H_{jt}^{o}) = \gamma g_{h,jt+1} (1 + g_{h,jt+1})^{2} \Phi_{jt+1} (H_{jt} - H_{jt-1})$$

where the last equality uses the envelope condition and the assumption that $H_{jt}^o = H_{jt-1}$. We can now define the labor wedge as the ratio of firm-specific wages over their marginal revenue product of labor:

$$\Lambda_{jt} \equiv \frac{\Phi_{jt}}{MRPH_{jt}} = \frac{\kappa\mu + (1-\kappa)\lambda}{\Gamma_{jt}}$$
(15)

where the wedge induced by labor adjustment costs (Γ_{jt}) can be written as:

$$\Gamma_{jt} \equiv \left(1 + \frac{\gamma}{2}g_{h,jt}^2\right) + (1 - \kappa)\lambda\left(\gamma g_{h,jt}(1 + g_{h,jt}) - \beta\gamma g_{h,jt+1}(1 + g_{h,jt+1})^2(1 + g_{\Phi,jt+1})\right) \\ - \kappa\gamma\beta g_{h,jt+1}(1 + g_{h,jt+1})^2(1 + g_{\Phi,jt+1})g_{h,jt}$$

The presence of adjustment costs imply that the labor wedge depends on current and future employment growth rates.

How will the labor wedge evolve over the firm's lifecycle? Suppose all firms begin their life with one employee. Due to convex adjustment costs, firms cannot instantaneously achieve their optimal size, but must grow towards it gradually. I now solve the firm's problem backwards from time \bar{T}_j , defined as the time at which firm *j* achieves its optimal steady state size. With this understanding, I now drop the subscript *j* for brevity. At time $t = \bar{T}$, firm size is such that $H_{t+1} = H_t = H_t^o$; firms no longer grow, and the labor wedge becomes $\bar{\Lambda} = \kappa \mu + (1 - \kappa)\lambda$. At time $\bar{T} - 1$, just before the firm achieves its optimal size, the labor wedge is:

$$\Lambda_{\bar{T}-1} = \frac{\Phi_{\bar{T}-1}}{\mu^{-1}\Omega^{\frac{\sigma-1}{\sigma}}H_{\bar{T}-1}^{-\frac{1}{\sigma}}}$$

Comparing time \overline{T} with $\overline{T} - 1$, we have that:

$$\frac{\Lambda_{\bar{T}}}{\Lambda_{\bar{T}-1}} = \frac{\Phi_{\bar{T}} H_{\bar{T}}^{\frac{1}{\sigma}}}{\Phi_{\bar{T}-1} H_{\bar{T}-1}^{\frac{1}{\sigma}}} = \left(\frac{H_{\bar{T}}}{H_{\bar{T}-1}}\right)^{\frac{1}{\eta}+\frac{1}{\sigma}} \ge 1 \implies \Lambda_{\bar{T}} \ge \Lambda_{\bar{T}-1}, \text{ since } H_{\bar{T}} \ge H_{\bar{T}-1}.$$

where the second equality uses the labor supply curve. Therefore, the labor wedge at steady state is larger than that one period before the firm achieves is steady state size. Extending this logic further back in time implies that the firm's labor wedge grows towards $\bar{\Lambda} = \kappa + (1 - \kappa)\lambda$ over time as firms grow towards their steady state size.

<u>Main implications</u>. This stylized model provides some guidance on how labor adjustment costs might bias the bargaining power and markdown estimates, and how one might assess the extent to which labor wedges reflect labor adjustment costs:

- (1.) All else equal, older firms have a higher labor wedge than younger firms, because they are closer to their optimal size. As firms grow closer to their optimal size, Λ converges to κµ + (1 – κ)λ from below. The extent to which firms' labor wedges rise as they age, conditional on their latent type, therefore contains information about the magnitude of labor adjustment costs in measured labor wedges.
- (2.) Similarly, all else equal, growing firms have a lower labor wedge than nongrowing firms (Λ_{jt} further below 1), because they are further from their optimal size.
- (3.) The above stylized example abstracts from variable markups, although my estimation of worker bargaining power relies on the presence of variable markups. To the extent that firms charge higher markups as they grow larger (towards their optimal size), there will be a positive correlation between markups and measured labor wedges, implying that $CV(\tilde{\mu}_{jt}, (1/\Gamma_{jt})) > 0$. Therefore, not accounting for labor adjustment costs could lead to an *upward biased* κ estimate.
- (4.) The presence of labor adjustment costs implies that one would *overstate* firm monopsony power λ for a given labor wedge Λ, markup μ, and bargaining power κ.

Incorporating labor adjustment costs into my baseline model. The previous model is stylized. It assumes that markups (μ) and markdowns (λ) are constant and homogenous, and that labor is the only factor input. I now incorporate labor adjustment costs into my structural framework in Section 2 for guidance on: (i) how to assess the extent to which labor wedges reflect labor adjustment costs; (ii) how to account for

labor adjustment costs when estimating bargaining power. I use the same quadratic labor adjustment cost function as the one above. The firm's state variables are now $Z_{jt} = \{\Omega_{jt}, D_{jt}, A_{jt}, H_{jt-1}, K_{jt}\}$, where capital K_{jt} is pre-determined.

Labor wedge. With labor adjustment costs, the labor wedge in my model (see equation (3)) becomes:

$$\Lambda_{jt} = \frac{\kappa \tilde{\mu}_{jt} + (1 - \kappa)\lambda_{jt}}{\Gamma_{jt}}$$
(16)

where the wedge induced by labor adjustment costs is:

$$\Gamma_{jt} \equiv \left(1 + \frac{\gamma}{2} g_{h,jt}^2 \right) + (1 - \kappa) \lambda \left(\gamma g_{h,jt} (1 + g_{h,jt}) - \beta \gamma E_t \left[g_{h,jt+1} (1 + g_{h,jt+1})^2 (1 + g_{\Phi,jt+1}) \right] \right)$$
$$- \kappa \gamma \beta E_t \left[g_{h,jt+1} (1 + g_{h,jt+1})^2 (1 + g_{\Phi,jt+1}) \right] g_{h,jt}$$

and $\tilde{\mu}_{jt} \equiv \left(1 - \frac{P_{m,t}M_{jt} + P_{k,t}K_{jt}}{P_{jt}Y_{jt}}\right) \frac{\mu_{jt}}{\alpha_{h,jt}}$ are product market rents.

Estimating equation for bargaining power. To assess whether labor adjustment costs might significantly bias my bargaining power estimates downwards, I derive the estimating equation for worker bargaining power than accounts for labor adjustment costs:

$$1 + \log \Lambda_{jt} \approx \kappa \tilde{\mu}_{jt} + (1 - \kappa) \lambda_{jt} - \log \Gamma_{jt}$$
(17)

Compared to the original estimating equation (3), this new equation includes the "adjustment cost wedge" Γ_{jt} . The derivation of equation (17) relies on the following firstorder approximations:

$$\log \left[\kappa \tilde{\mu}_{jt} + (1-\kappa)\lambda_{jt}\right] \approx \kappa \tilde{\mu}_{jt} + (1-\kappa)\lambda_{jt} - 1$$
, and

$$\log \Gamma_{jt} \approx \frac{\gamma}{2} g_{h,jt}^{2} + (1-\kappa) \lambda \left(\gamma g_{h,jt} (1+g_{h,jt}) - \beta \gamma E \left[g_{h,jt+1} (1+g_{h,jt+1})^{2} (1+g_{\Phi,jt+1}) \middle| Z_{jt} \right] \right) \\ - \kappa \gamma \beta E \left[g_{h,jt+1} (1+g_{h,jt+1})^{2} (1+g_{\Phi,jt+1}) \middle| Z_{jt} \right] g_{h,jt}$$

I discuss next how I use equation (17) to assess the magnitude of labor adjustment costs and estimate worker bargaining power.

Assessing the extent to which the labor wedge (Λ_{jt}) reflects labor adjustment costs (Γ_{jt}) . In light of the discussion above on how labor wedges are expected to differ by firm age

and employment growth rates (see main implications (1.) and (2.) of the deterministic model), I now compare the estimated labor wedges $\log \hat{\Lambda}_{jt}$ along these two observed firm characteristics. Motivated by equation (17), I run the following regression:

$$\log \hat{\Lambda}_{it} = \text{Category}_{it} + \mathcal{X}' \mathcal{B} + \varepsilon_{it}$$

where Category_{*jt*} refers categories of firm age and employment growth rates to which firm *j* belongs at time *t*. Firm age is defined by ten groups of 5-year increments, starting at the age group between 0 to 5 years old. Employment growth rate categories are defined in terms of deciles. The vector of controls \mathcal{X} includes product market rents $\hat{\mu}_{jt}$, firm size as measured by employment and wage bills (implied by $\lambda_{jt} = \lambda(H_{jt}, \mathcal{H}^{-1}(H_{jt}, \Phi_{jt}H_{jt})))$, firm fixed effects, and sector×year fixed effects. I restrict the sample to a balanced panel of firms that appeared in every year between 2009 and 2016, and on average experienced a positive growth rate in employment in this time frame. I present the findings in Figures 17 and 18. Each figure contains four panels. In panel (a), \mathcal{X} includes firm size controls and sector×year fixed effects. In panel (b), \mathcal{X} additionally includes product market rents. In panel (c), \mathcal{X} additionally includes firm fixed effects. In panel (d), \mathcal{X} additionally includes both product market rents and firm fixed effects.

Figure 17 presents how labor wedges vary depending on the age of the firm. Panel (a) shows that labor wedges increase with firm age. Panels (b), (c), and (d) show that controlling for product market rents or looking at within-firm variation significantly flattens the profile of labor wedges across firm age. A similar pattern holds when looking at employment growth rates, as Figure 18 shows. When I control for product market rents or look at within-firm variation (or both), the profile of labor wedges along firm employment growth rates largely flattens. Overall, these findings suggest that while labor adjustment costs may be present, they seem unlikely to be the main determinant of labor wedges or to significantly affect the bargaining power estimates. Nevertheless, I now attempt to account for labor adjustment costs when estimating worker bargaining power.



Figure 17: Labor wedges by firm age.

Note: Panel (a) includes firm size controls and sector × year fixed effects. Panel (b) additionally includes product market rents. Panel (c) additionally includes firm fixed effects, but not product market rents. Panel (d) additionally includes both product market rents and firm fixed effects. Vertical bars are 95% confidence intervals.



Figure 18: Labor wedges by employment growth rate.

Note: Panel (a) includes firm size controls and sector × year fixed effects. Panel (b) additionally includes product market rents. Panel (c) additionally includes firm fixed effects, but not product market rents. Panel (d) additionally includes both product market rents and firm fixed effects. Vertical bars are 95% confidence intervals.

Accounting for labor adjustment costs when estimating worker bargaining power (κ).

I now adapt my approach (proposed in Section 3) to estimating bargaining power to account for labor adjustment costs using equation (17). The control function approach taken to address unobserved amenities when estimating κ is unchanged—the monopsony markdowns can still be written as a function of firm size in employment and wage bills $\lambda(H_{jt}, \mathcal{H}^{-1}(H_{jt}, \Phi_{jt}H_{jt}))$. However, the presence of the labor adjustment wedge (Γ_{jt}) introduces the following component into equation (17):

$$E\left[\underbrace{g_{h,jt+1}(1+g_{h,jt+1})^2(1+g_{\Phi,jt+1})}_{\equiv g_{j,t+1}^{adj}}\Big|Z_{jt}\right]$$

Therefore, the estimation procedure is now complicated by the presence of *expected* future employment and wage growth rates ($g_{h,jt+1}$ and $g_{\Phi,jt+1}$), which are unobserved.

To address unobserved future employment and wage growth rates, I replace the t + 1 growth rates expected at time t with the actual growth rates at t + 1 in the expectation term above. That is, I compute the realized version of the term in expectations above, $g_{j,t+1}^{adj} \equiv g_{h,jt+1}(1+g_{h,jt+1})^2(1+g_{\Phi,jt+1})$. I then note that, under the assumption of rational expectations, expected growth rates ($E[g_{j,t+1}^{adj}|Z_{jt}]$) are on average equal to realized growth rates ($g_{j,t+1}^{adj}$). Under this assumption, deviations of actual growth rates from their expectations can be thought of as random expectation errors, or "measurement error". This gives a *lower bound* estimate for κ due to attenuation bias.

In an alternative approach to addressing the unobserved expected future growth rates, I project realized growth rates at time t + 1, $g_{j,t+1}^{adj}$, on the firm's state variables Z_{jt} to compute the expected future growth rates at time t, $E[g_{j,t+1}^{adj}|Z_{jt}]$. That is, I use the predicted future growth rates at time t ($\hat{g}_{j,t+1}^{adj}$) in place of the expected growth rates ($E[g_{j,t+1}^{adj}|Z_{jt}]$) in equation (17). However, among the state variables, quality D_{jt} and amenities A_{jt} are unobserved. When computing the predicted future growth rates $\hat{g}_{j,t+1}^{adj}$, I can only project $g_{j,t+1}^{adj}$ on $\hat{\Omega}_{jt+1}$, H_{jt} , and K_{jt+1} . Therefore, to the extent that a higher current quality or amenity raises expected future growth rates, these will be omitted variables that upward bias the estimated κ . This gives an *upper bound* estimate for κ .

Table 20 shows the estimates of worker bargaining power under regression specifications that use realized future growth rates (columns 1 and 2), and those that use

	Labor wedges (Λ_{it})						
	(1)	(2)	(3)	(4)			
Product market rents ($\tilde{\mu}_{jt}$)	0.159	0.100	0.159	0.100			
	(0.005)	(0.007)	(0.005)	(0.007)			
Controls for $E[g_{i,t+1}^{adj} Z_{jt}]$	Realized	Realized	Predicted	Predicted			
Sector × Year FE	\checkmark	\checkmark	\checkmark	\checkmark			
Firm FE		\checkmark		\checkmark			
Observations	68,623	65,952	68,623	65,952			

Table 20: Estimates of κ , controlling for quadratic labor adjustment costs.

This table presents estimates of worker bargaining power (κ) that accounts for quadratic labor adjustment costs. I control for expected future growth rates, $E[g_{j,t+1}^{adj}|Z_{jt}]$, induced by the presence of labor adjustment costs, in two ways, using realized $g_{j,t+1}^{adj}$ (denoted as "Realized") and using predicted $g_{j,t+1}^{adj}$ (denoted as "Predicted"). Bootstrapped standard errors in parentheses. All regressions in this table include fourth-order polynomials of employment and wage bills to address unobserved amenities, following the control function approach for monopsony markdowns detailed in Section 3.3.

predicted future growth rates (columns 3 and 4). Column 2 shows that the point estimate for κ of 0.100 is similar to my baseline estimate in column 4 of Table 2 of 0.125. Column 4 in Table 20 shows that the estimates are similar when I use predicted future growth rates in place of realized growth rates.

B.12 Estimating labor wedges using hiring wages only

A caveat for the results presented in Section 5 is that firm wage premia and labor wedges are estimated for all workers; both new hires and incumbent workers. As discussed in Section 3, it may be important to allow the wages of incumbent workers to be determined separately from those of new hires. In this section, I estimate labor wedges using hiring wages only, following Di Addario et al. (2020). I do not take a stance on the wage-setting protocol for incumbent workers.

To repeat the main estimation exercise of this paper, I first k-means cluster firms into groups using only hiring wages (W^n) and estimate firm wage premia (ϕ^n). To estimate production functions taking into account differences in worker efficiency, I compute the average worker efficiency at each firm using the following relationship: $W_{it}^n = \bar{E}_{it}^n \Phi_{it}^n$. The rest of the estimation routine is as described in Section 3.

Once production functions and price-cost markups are estimated, labor wedges are measured as follows:

$$\Lambda_{jt}^n = \frac{W_{jt}^n L_{jt}}{P_{jt} Y_{jt}} \cdot \mu_{jt} \cdot \alpha_{h,jt}^{-1}$$

 Λ^n represents the labor wedges for new hires.

Table 21 below shows that the estimated price-cost markups are similar to those estimated using all workers in Table 1 in Section 5. However, the estimated labor wedges are lower in this case: new hires are paid a lower share of their marginal revenue product than incumbent workers. This is consistent with the findings of Kline et al. (2019), who show that patent-induced labor productivity shocks pass through to incumbent workers' wages, but not the wages of new hires.

	Mean	Median	25 th Pct	75 th Pct	Var	Var (i)	Var (ii)
Firm wage premium	2.97	2.97	2.91	3.02	0.01	0.00	0.02
Labor wedges	0.68	0.71	0.54	0.89	0.16	0.09	0.11
Markups	1.33	1.33	1.14	1.60	0.09	0.05	0.03
# firms				14,342			

Table 21: Summary statistics for estimated firm wage premia and labor wedges in 2016 (using hiring wages only).

This table reports the summary statistics for estimated firm wage premia and labor wedges using hiring wages only (Di Addario et al., 2020). Variances are reported for the logarithms of those variables. The column Var (i) reports the variances corrected for measurement error following Krueger and Summers (1988) and Kline et al. (2020), while the column Var (ii) reports the variances for firm-groups. Markups and labor wedges are winsorized by 2%.

B.13 Bootstrapped standard errors

In Section 3.3, I use the labor wedge equation (3) to estimate worker bargaining power (κ). Implementing this equation requires measuring labor wedges (λ_j) and product market rents ($\tilde{\mu}_j$), which requires estimating production functions. Tables (2), (18), and (20) therefore report bootstrapped standard errors for the κ estimates. I now describe the bootstrapping procedure I implement to compute those standard errors:

- i. For each two-digit French manufacturing sector, I randomly sample (with replacement) 90% of firms. For each sampled firm, I extract their entire panel in the time span that I observe them (2009-2016).
- ii. I estimate the production function for each two-digit sector using the sampled firms and compute the labor wedges (λ_i) and product market rents ($\tilde{\mu}_i$).
- iii. I then implement the labor wedge equation (3) on only the sampled firms to estimate κ , then compute their λ_i and $\tilde{\mu}_i$.

I repeat steps (i.) to (iii.) 500 times and compute the bootstrapped standard errors.

C Appendix: Theory

C.1 Discussion: Individual outside options

In my model in Section 2 (and Section 6), I explore the role of worker bargaining power when without which firms would behave as monopsonist wage-posters.³¹ Since a monopsonist chooses a wage on the labor supply curve that makes the marginal worker indifferent between the firm and its next best competitor, the marginal worker and all inframarginal ones are individually willing to accept the monopsonist's wage offer. What, then, might grant those workers the power to ask for wages above the monopsony wage?

One approach could be to follow the model of on-the-job search and individual bargaining of Cahuc et al. (2006), allowing firms to renegotiate wages with individual employees on a case-by-case basis, depending on the specific outside offers their employees receive. This would introduce worker-specific outside options into the wage bargaining problem.

Incorporating on-the-job search and an individual wage bargaining protocol that accounts for heterogeneous outside options is appealing: it would allow researchers to study the implications of specific labor market arrangements (such as non-compete agreements) that may impede employer-to-employer transitions or individual workers' ability to trigger wage renegotiations with their employers.

However, incorporating these interesting aspects is also very challenging in the type of environment in Section 2: firms have labor market power and product market power. The presence of product market power implies that firms have decreasing-returns-to-scale revenue functions, which brings a set of complications with modeling labor market power from an on-the-job search and individual bargaining perspective.

One key challenge, as Mortensen and Vishwanath (1991) point out, is that the presence of decreasing returns in combination with on-the-job search can lead to no equilibrium or multiple equilibria.³² Therefore, I abstract from on-the-job search as a microfoundation for labor supply curves—and hence labor market power—in my model.

³¹Here, I use the term 'monopsonist' in the sense that firms have wage-setting power, instead of the literal sense of having only one buyer.

³²Bilal and Lhuillier (2025) develop a recent approach to solving this problem, which requires that firm employment monotonically increases with their productivity. However, the presence of non-wage amenities in my model violates this monotonicity requirement.

The presence of decreasing returns also introduces challenges with modelling bilateral bargaining between an individual worker and a firm: the marginal product of all workers in a firm now depends on the bargaining outcome with one worker. If bargaining breaks down with that worker, the marginal product (and hence, wages) of their coworkers increase. This makes the bargaining problem more complex than with constant returns (Stole and Zwiebel, 1996). This individual bargaining problem becomes yet more complex if one allows the possibility that firms may derive market power from their size (i.e., firms are non-atomistic), as in the more specific version of my model in Section 6 or other recent non-bargaining models (Atkeson and Burstein, 2008; Berger et al., 2022). In such environments, the firm's outside option not only has to account for how an impasse in bargaining with one worker would affect the wages of its other employees, but also how its product and labor market competitors would respond to those wage changes, which in turn affects its own employment and production decisions, as well as wage bargaining problem.

The modeling approach taken in this paper—that workers *collectively* bargain with the employer-avoids these challenges and is consistent with French labor market institutions. This distinction between individual and collective bargaining is important, as it is determines both the workers' and firm's relevant outside options in the wage bargaining process. In my model, I assume that workers can collectively threaten to go on a one-period strike by temporarily halting production, thereby penalizing their employer with zero profits while incurring a zero payoff themselves over the duration of the strike (since no output is produced).³³ Since the model is static and bargaining occurs every period, the bargaining problem with the same workers is repeated next period. This collective bargaining enables workers to demand a wage above the monopsony wage, with the extent depending on workers' bargaining strength. This approach is similar to recent work on labor unions. For example, Taschereau-Dumouchel (2020) models unionized firms as paying higher wages than non-unionized firms because of workers' ability to collectively quit into unemployment if no agreement is reached. Nevertheless, the more conventional approach to modeling wage bargaining is between an individual worker and a firm (i.e., individual bargaining).

³³While going on strike is not the same as quitting into unemployment, this zero payoff assumption is consistent with the lack of wage response from large increases in unemployment insurance levels (Jäger et al., 2020), and the lack of response of reservation wages to changes in the potential unemployment benefit duration (Le Barbanchon, Rathelot, and Roulet, 2019).

Firm-level collective wage bargaining is consistent with French labor market institutions. French labor market institutions stipulate that firms with at least 50 employees must bargain annually with a union representative who represents workers at that firm (see Appendix A.1). For such firms, the presence of at least one union representative is a binding legal requirement. Firm-level collective wage bargaining is not only a French institution. As Bhuller et al. (2022) point out, the vast majority of workers in Europe are covered by collective bargaining agreements (see their Figure 1) and collective bargaining at the firm-level has become a prevalent approach to wage determination in OECD countries (see their Figure 2). However, as Bhuller et al (2022) also note, not all countries allow a full-scale strike in wage negotiations.³⁴

By modeling bargaining power as arising from workers' ability to take collective action, the model is more closely tied to the discussion on whether raising worker bargaining power through strengthening pro-worker institutions could lead to welfare gains (see, e.g. Stansbury and Summers (2020)), where topics such as collective bargaining agreements, unionization, profit-sharing plans, and workers on corporate boards are a central focus. However, as discussed above, the omission of on-the-job search and individual-level bargaining also means that my model is not well-suited for analyzing the implications of policies that might strengthen individual bargaining positions, such as banning non-compete agreements.

An interesting avenue for future work could be to incorporate a wage renegotiation game (such as Cahuc et al. (2006)) between an individual worker and their employer to compare the relative importance of individual and collective wage bargaining in determining the size of firms' labor market power. In such an extension, the collectively bargained wage can be thought of as a base wage at a given employer, and the realized wage an individual worker receives may deviate upwards depending on the outside offers they receive. However, attempts at such an extension would require addressing the tractability challenges described by Mortensen and Vishwanath (1991) and Stole and Zwiebel (1996).

³⁴For example, in Norway and Sweden, a full-scale strike is not allowed, though workers may reduce production substantially without halting it altogether.

C.2 Firm wage premia with outside wages

As before, workers bargain collectively with their employer j and bargaining is efficient: workers and firms jointly choose wages, prices, materials, and capital to maximize total rents, taking into account the product demand curve and labor supply curve. Firms have an outside option of zero profits. The firm-specific labor supply curve is now $H_j = \mathcal{H}(\Phi_j - \Phi^o, A_j)$; workers do not supply labor unless firms pay at least the exogenous outside wage Φ^o . Workers and firms maximize the following Nash product:

$$\max_{\Phi_j, P_j, M_j, K_j} \left((\Phi_j - \Phi^o) H_j \right)^{\kappa} \left(\Pi_j \right)^{1-\kappa}$$

subject to $H_j = \mathcal{H}(\Phi_j - \Phi^o, A_j)$, $Y_j = \mathcal{G}(P_j, D_j)$, and $Y_j = \Omega_j F(K_j, M_j, H_j)$. The firm's profit is $\Pi_j = P_j Y_j - \Phi_j H_j - P_m M_j - P_k K_j$. The firm-specific wage premium is:

$$\Phi_{j} = \Phi^{o} + \kappa \left(\underbrace{\frac{P_{j}Y_{j} - P_{m}M_{j} - P_{k}K_{j}}{H_{j}}}_{\text{Total rents per effective labor}}\right) + (1 - \kappa) \underbrace{\left[\lambda_{j}(MRPH_{j} - \Phi^{o})\right]}_{\text{Monopsony wages}}$$
(18)

Equation (18) shows that the firm wage premium is again a weighted average of a pure bargaining outcome and a pure monopsony outcome. When $\kappa = 0$, workers receive a mark down of monopsony rents in addition to the outside wage. The markdown λ_j is determined by the firm-specific labor supply elasticity. In this case, workers do not receive rents from the firm's product market power. When $\kappa = 1$, workers receive the total amount of rents generated by firms' labor and product market power.

Firm wage premia can be written in exactly the same form as in equation (2). The labor wedge Λ_i in this case is:

$$\Lambda_j = \frac{1}{1 - (1 - \kappa)(1 - \lambda_j)\frac{\Phi^o}{\Phi_j}} \left[\kappa \left(1 - \frac{\alpha_{m,j} + \alpha_{k,j}}{\mu_j} \right) \frac{\mu_j}{\alpha_{h,j}} + (1 - \kappa)\lambda_j \right]$$
(19)

Labor wages are therefore higher when workers have outside wages, all else equal.

Estimating workers' bargaining power. In Section 3, I proposed a control function approach to address unobserved variation in amenities when estimating bargaining power. The estimating equation assumed that outside wages are zero. I now adjust the estimating equation to account for positive outside wages. The adjusted estimating
equation is:

$$\tilde{\Lambda}_{jst} = \kappa \mathcal{M}_{jst} + (1 - \kappa) \lambda(H_{jst}, \mathcal{A}(H_{jst}, \Phi_{jst}H_{jst})) + \tilde{\epsilon}_{jst}$$

where $\tilde{\Lambda}_{jst} \equiv \frac{\Phi_{jst} - \Phi_{jst}^o}{MRPH_{jst} - \Phi_{jst}^o}$, $\mathcal{M}_{jst} \equiv \frac{\left(1 - \frac{\alpha_{m,j} + \alpha_{k,j}}{\mu_j}\right)\frac{\mu_j}{\alpha_{h,j}}MRPH_{jst} - \Phi_{jst}^o}{MRPH_{jst} - \Phi_{jst}^o}$, and $\tilde{\epsilon}$ is a residual capturing measurement error. As in Section 3, I approximate the markdown function $\lambda(.)$ with a 4th-order polynomial in employment and wage bills. The markdowns are then computed using the estimated bargaining parameter: $\lambda_{jst} = \frac{1}{1-\hat{\kappa}} \left(\tilde{\Lambda}_{jst} - \hat{\kappa}\mathcal{M}_{jst} - \hat{\epsilon}_{jst}\right)$.

Implementing the estimation procedure requires measuring the outside wage. One option is to set Φ^o equal to the level of unemployment benefits. However, Jäger et al. (2020) show that even large reforms to unemployment benefit levels in Austria did not materially affect wages. Similarly, Le Barbanchon et al. (2019) show that extending the potential unemployment benefit duration in France did not affect reservation wages reported by jobseekers. Because the model implies that firms paying below the outside wage will not hire any workers, I set the outside wage to the minimum observed wage.

The estimated workers' bargaining power using this measure of outside wages is reported in Table 22. These estimates are smaller than the ones that assume zero outside wages in Table 2. I then compute the implied monopsony wage markdown, taking the estimated bargaining power in column (4) of Table 22. The wage markdowns at the median firm is 0.45 under $\Phi^{o} = 0$ in Table 3. The markdowns at the median firm computed with the current measure of Φ^{o} is 0.47.

	Labor wedge (Λ_{jt})					
	(1)	(2)	(3)	(4)	(5)	(6)
Product market	0.088	0.051	0.088	0.050	0.093	0.071
rents ($\tilde{\mu}_{jt}$)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.014)
Sector×year fixed effects	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark
Firm fixed effects		\checkmark		\checkmark		\checkmark
Control for $\lambda(.)$			\checkmark	\checkmark	\checkmark	\checkmark
IV for measurement error					\checkmark	\checkmark
Observations	101.584	98.624	101.584	98.624	78,434	75,788

Table 22: Estimated worker bargaining power (when workers have outside wages).

This table reports the estimated workers' bargaining power parameter when workers have outside wages. Even columns controls for firm fixed effects. Columns (3) through (6) controls for differences in monopsony markdowns reflecting differences in amenities. Columns (5) and (6) instruments measured product market rent $\tilde{\mu}_{jt}$ with its lag to address classical measurement error. Bootstrapped standard errors are reported in parentheses.

Table 23: The distribution of monopsony markdowns (with non-zero outside wages).

Summary statistics	Mean	Median	25 th Pct	75 th Pct	Var	Var (i)	Var (ii)
Markdowns (λ_{jt})	0.50	0.47	0.36	0.60	0.16	0.16	0.07

This table reports the summary statistics in 2016 for the estimated monopsony markdowns when workers have outside wages. Variances are reported for *log* monopsony markdowns. The column Var (i) reports the variances corrected for measurement error following Krueger and Summers (1988) and Kline et al. (2020), while the column Var (ii) reports the variances for firm-groups. Each variable is winsorized by 2%.

C.3 Firm wage premia with Stole & Zwiebel (1996) bargaining

Consider a setting in which firms bargain individually with their employees. Workers supply one unit of labor inelastically. Firms produce goods with only labor inputs. The labor market is characterized by search frictions. Firms post vacancies V_j subject to a convex vacancy posting cost $c(V_j)$. The product market is imperfectly competitive; firms face the following product demand curve $Y_j = P_j^{-\sigma}$. The production function is $Y_j = \Omega_j H_j^{\alpha_h}$. There are no idiosyncratic shocks to firms. Firms steady state size is $H_j = \frac{q}{\delta}V_j$, where q is the job-filling rate and $\overline{\delta}$ is the separation rate. Firms post vacancies to maximize profits:

$$\max_{V_j} P_j Y_j - \Phi_j H_j - c(V_j)$$

subject to the product demand curve and the firm size constraint. The solution to the Stole and Zwiebel (1996) bargaining problem is then:

$$(1-\kappa)(\Phi_j - \Phi_j^o) = \kappa(MRPH_j - \frac{\partial \Phi_j}{\partial H_j}H_j - \Phi_j)$$

Solving this differential equation yields the following firm wage premium equation:

$$\Phi_j = \left(\frac{\kappa\mu\alpha_h^{-1}}{(1-\kappa)\mu\alpha_h^{-1} + \kappa}\right) MRPH_j + (1-\kappa)\Phi^o$$
(20)

where markups are constant $\mu = \frac{\sigma}{\sigma-1}$ and the marginal revenue product of labor is $MRPH_j = \mu^{-1} \alpha_h \frac{P_j Y_j}{H_i}$.

Comparison with collective (efficient) bargaining. When $\alpha_k = \alpha_m = 0$ and labor supply is completely inelastic (as is the case when deriving equation (20) under Stole-

Zwiebel bargaining), the firm wage premium equation under collective bargaining (18) becomes:

$$\Phi_j = \kappa \mu \alpha_h^{-1} M R P H_j + (1 - \kappa) \Phi^o$$
⁽²¹⁾

Comparing equation (21) with equation (20) shows that, under collective (efficient) bargaining, workers are able to extract a higher share of *MRPH* as wages than under individual bargaining, since $\frac{1}{(1-\kappa)\mu\alpha_h^{-1}+\kappa} \leq 1$.

C.4 The Social Planner's problem

The planner maximizes: :

$$\max_{H_j, K_j M_j} C - \frac{H^{1+\varphi}}{1+\varphi}$$

subject to:

$$C + K + M = Y, \qquad Y = \left[\int_{0}^{1} \tilde{D}_{g}^{\frac{1}{\theta}} Y_{g}^{\frac{\theta-1}{\theta}} dg\right]^{\frac{\theta}{\theta-1}}, \qquad Y_{g} = \left(\sum_{j \in g} \tilde{D}_{jg}^{\frac{1}{\sigma}} Y_{jg}^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}}$$
$$K = \int_{g} \sum_{j \in g} K_{jg} dg, \qquad M = \int_{g} \sum_{j \in g} M_{jg} dg, \qquad Y_{j} = \Omega_{j} K_{j}^{\alpha_{k}} M_{j}^{\alpha_{m}} H_{j}^{\alpha_{h}}$$
$$H = \left[\int_{0}^{1} \tilde{A}_{s}^{-\frac{1}{\nu}} H_{s}^{\frac{\nu+1}{\nu}} ds\right]^{\frac{\nu}{1+\nu}}, \qquad H_{s} = \left(\sum_{j} \tilde{A}_{js}^{-\frac{1}{\eta}} H_{js}^{\frac{\eta+1}{\eta}}\right)^{\frac{\eta}{\eta+1}}$$

The planner's aggregate and firm-level labor allocations are characterized by:

$$H^{\varphi} = \underbrace{\alpha_{h} \frac{Y}{H}}_{MPL} \quad \text{and} \quad H^{\varphi} \left(\frac{H_{js}}{H_{s}}\right)^{\frac{1}{\eta}} \left(\frac{H_{s}}{H}\right)^{\frac{1}{\nu}} = \underbrace{\alpha_{h} \frac{Y_{jg}}{H_{js}}}_{MPL_{js}} \tilde{D}_{jg}^{\frac{1}{\theta}} \tilde{D}_{jg}^{\frac{1}{\varphi}} \left(\frac{Y_{jg}}{Y_{g}}\right)^{-\frac{1}{\varphi}} \left(\frac{Y_{g}}{Y}\right)^{-\frac{1}{\theta}}$$

C.5 Market power and misallocation: subsidizing firm size

In Section 6, I assess the importance of variable markups, markdowns, and worker bargaining power for wages and welfare. To implement either the social planner's equilibrium or an efficient allocation of factor inputs across firms, I derive a set of subsidies that accomplish this goal, following Edmond et al. (2023). These subsidies are meant as a tool to quantify the aggregate welfare effects of markups or markdowns, and to quantify how far raising worker bargaining power can go in raising aggregate welfare. My model does not speak to the distributional welfare effects of these subsidies as I do not model household heterogeneity. In what follows, I start by deriving the output and labor subsidies that offset firm market power when workers have no bargaining power ($\kappa = 0$). I then show how the subsidies are affected when workers have some bargaining power. I assume that these subsidies are financed by a lump sum tax on the household.

Case 1: No worker bargaining power. Suppose workers have no bargaining power ($\kappa = 0$). For a given set of size-based subsidies $T_y(Y_{jg})$ and $T_h(H_{js})$, the firm now maximizes:

$$\max_{\Phi_{js}, Y_{jg}, K_j, M_j} \Pi_j^{gross} = P_{jg} Y_{jg} - \Phi_{js} H_{js} - P_k K_j - P_{m,j} M_j + T_y (Y_{jg}) + T_h (H_{js})$$

subject to the production function $Y_j = \Omega_j K_j^{\alpha_k} M_j^{\alpha_m} H_j^{\alpha_h}$, labor supply curve (8), and product demand curve (9).

Equalizing or removing markups. The first-order conditions for capital and materials are:

Capital FOC:
$$\left(P_{jg} + Y_{jg} \frac{\partial P_{jg}}{\partial Y_{jg}} + T'_{y}(Y_{jg}) \right) \alpha_{k} \frac{Y_{jg}}{K_{j}} = P_{k}$$

Materials FOC:
$$\left(P_{jg} + Y_{jg} \frac{\partial P_{jg}}{\partial Y_{jg}} + T'_{y}(Y_{jg}) \right) \alpha_{m} \frac{Y_{jg}}{M_{j}} = P_{m,j}$$

The term $Y_{jg} \frac{\partial P_{jg}}{\partial Y_{jg}}$ captures the distortion induced by the firm's price-setting power. As in Section 6, let ρ_{jg} be the price elasticity of demand. A subsidy schedule that satisfies:

$$T'_{y}(Y_{jg}) = -Y_{jg} \frac{\partial P_{jg}}{\partial Y_{jg}} = \frac{1}{\rho_{jg}} P_{jg}$$

offsets the firm's price-setting power, inducing the firm to behave as a pricetaker in the product market. Alternatively, a subsidy schedule that satisfies:

$$T'_{y}(Y_{jg}) = \left(\frac{1}{\rho_{jg}} - \frac{1}{\bar{\sigma}}\right) P_{jg}$$
(22)

induces firms to charge a constant markup $\bar{\mu} = \frac{\bar{\sigma}}{\bar{\sigma}-1}$. Any level of constant markups can be implemented by an appropriate choice of the policy parameter $\bar{\sigma}$. For example, $\bar{\sigma} \to \infty$ induces firms to behave as pricetakers. The output subsidy schedule is then

 $T_y(Y_{jg}) = \int T'_y(Y_{jg}) dY_{jg}$. These results have been shown by Edmond et al. (2023).

Equalizing or removing monopsony markdowns. I now write the first-order condition for labor to show that an appropriately chosen labor subsidy schedule can offset firm's labor market power:

Labor FOC:
$$\left(P_{jg} + Y_{jg}\frac{\partial P_{jg}}{\partial Y_{jg}} + T'_{y}(Y_{jg})\right)\alpha_{h}\frac{Y_{jg}}{H_{js}} = \Phi_{js} + H_{js}\left(\frac{\partial H_{js}}{\partial \Phi_{js}}\right)^{-1} - T'_{h}(H_{js})$$

The term $H_{js} \left(\frac{\partial H_{js}}{\partial \Phi_{js}}\right)^{-1}$ captures firms' wage-setting power and distorts labor allocations away from the social planner's allocation. A labor subsidy schedule that satisfies:

$$T'_{h}(H_{js}) = H_{js} \left(\frac{\partial H_{js}}{\partial \Phi_{js}}\right)^{-1} = \frac{1}{\xi_{js}} \Phi_{js}$$

offsets firms labor market power, where ξ_{js} is the labor supply elasticity. Alternatively, a labor subsidy schedule that satisfies:

$$T'_{h}(H_{js}) = \left(\frac{1}{\xi_{js}} - \frac{1}{\bar{\eta}}\right) \Phi_{js}$$
(23)

induces firms to mark wages down by a constant $\bar{\lambda} = \frac{\bar{\eta}}{1+\bar{\eta}}$. That is, $\bar{\eta}$ is a policy parameter than can be chosen to implement a specific level of the constant markdown $\bar{\lambda}$. The labor subsidy schedule is then $T_h(H_{js}) = \int T'_h(H_{js}) dH_{js}$. These results extend the results in Edmond et al. (2023), whose model does not feature imperfect labor market competition.

Case 2: Positive worker bargaining power. I now derive the output and labor subsidies that address firm product and labor market power when workers have some bargaining power $\kappa > 0$. For convenience, I restate the bargaining problem between workers and their employer:

$$\max_{\Phi_{js}, Y_{jg}, K_j, M_j} \left(\Phi_{js} H_{js} \right)^{\kappa} \left(\Pi_j^{gross} \right)^{1-\kappa}$$

subject to the production function $Y_j = \Omega_j K_j^{\alpha_k} M_j^{\alpha_m} H_j^{\alpha_h}$, labor supply curve (8), and product demand curve (9). The gross profits are $\Pi_j^{gross} = P_{jg}Y_{jg} - \Phi_{js}H_{js} - P_kK_j - P_kK_j$

 $P_{m,j}M_j + T_y(Y_{jg}) + T_h(H_{js}).$

Equalizing or removing markups. To equalize or remove *only* markup-induced distortions, one can use the same output subsidy schedule as in Case 1 (where $\kappa = 0$), because the first-order conditions for capital and materials are the same. However, when workers have bargaining power, workers will appropriate some of the profits due to the output subsidy. That is, workers benefit directly from the output subsidy, and not only through the effects of the subsidy on the firm's pricing behavior. Therefore, to isolate the welfare effects solely from equalizing markups, an appropriate set of labor *taxes* on the firm needs to be designed. To see this, consider the wage equation under output subsidies, but without any labor taxes:

$$\Phi_{js} = \kappa Q_{js}^{rent} + (1 - \kappa) \left[\left(\frac{\xi_{js}}{1 + \xi_{js}} \right) \left(1 - \frac{1}{\rho_{jg}} + \frac{T'_y(Y_{jg})}{P_{jg}} \right) \alpha_h \frac{P_{jg}Y_{jg}}{H_{js}} \right] + \kappa \frac{T_y(Y_{jg})}{H_{js}}$$
(24)

where $Q_{js}^{rent} \equiv \frac{P_{jg}Y_{jg} - P_kK_j - P_{m,j}M_j}{H_{js}}$. Equation (24) is almost identical to the original wage equation (1) where there are no output subsidies; the main differences are that (i) markups are now offset by the output subsidy $(T'_y(Y_{jg}))$, and (ii) workers capture some of the output subsidies to the firm because they have bargaining power ($\kappa \frac{T_y(Y_{jg})}{H_{js}}$). With labor taxes, this equation becomes:

$$\Phi_{js} = \kappa Q_{js}^{rent} + (1 - \kappa) \left[\left(\frac{\xi_{js}}{1 + \xi_{js}} \right) \left(1 - \frac{1}{\rho_{jg}} + \frac{T'_y(Y_{jg})}{P_{jg}} \right) \alpha_h \frac{P_{jg}Y_{jg}}{H_{js}} \right] + \kappa \left(\frac{T_y(Y_{jg})}{H_{js}} + \frac{T_h(H_{js})}{H_{js}} \right) + (1 - \kappa)T'_h(H_{js})$$
(25)

Equation (25) shows that to isolate the welfare effects of equalizing markups, without capturing the direct effect of the output subsidy through worker bargaining power—that is, without capturing $\kappa \frac{T_y(Y_{jg})}{H_{js}}$ —the appropriate labor tax must satisfy:

$$\kappa \left(\frac{T_y(Y_{jg})}{H_{js}} + \frac{T_h(H_{js})}{H_{js}}\right) + (1-\kappa)T'_h(H_{js}) = 0$$
(26)

Equalizing or removing monopsony markdowns. Suppose that a policymaker wishes only to equalize or remove markdowns and is not interested in markups ($T_y(Y_{js}) = 0$). Then, given equation (24), the labor subsidy schedule that achieves the policymaker's

goal of implementing a monopsony markdown of $\overline{\lambda}$ must satisfy:

$$(1-\kappa)T_h'(H_{js}) = (1-\kappa)\left(\frac{1}{\xi_{js}} - \frac{1}{\bar{\eta}}\right)\Phi_{js} - \kappa\left(1 + \frac{1}{\xi_{js}}\right)\frac{T_h(H_{js})}{H_{js}}$$
(27)

The first term on the right-hand side of the equality addresses firm monopsony power. When $\bar{\eta} \to \infty$, the subsidy removes monopsony markdowns entirely, although it does not remove the labor wedge Λ_{js} because workers are still able to appropriate profits from markups. The second term takes into account that workers are able to capture a part of the labor subsidy given to the firm. If workers have no bargaining power ($\kappa = 0$), then the labor subsidy reduces to the one in Case 1: $T'_h(H_{js}) = \left(\frac{1}{\xi_{js}} - \frac{1}{\bar{\eta}}\right) \Phi_{js}$.

Finally, if the policymaker also chooses output subsidies to equalize markups, then the labor subsidy that implements constant monopsony markdowns must satisfy:

$$(1-\kappa)T'_{h}(H_{js}) = (1-\kappa)\left(\frac{1}{\xi_{js}} - \frac{1}{\bar{\eta}}\right)\Phi_{js} - \kappa\lambda_{js}^{-1}\frac{T_{y}(Y_{jg}) + T_{h}(H_{js})}{H_{js}}$$
(28)

Welfare effects of the output and labor subsidies. The output and labor subsidies only affect the household's budget constraint by affecting firms' input demand decisions they do not otherwise directly affect the household's budget constraint. There are two reasons for this. First, because the output and labor subsidies are financed by a lump sum tax on the representative household, the subsidies do not distort household consumption and labor supply decisions. Second, because firm profits (inclusive of the subsidies) are rebated to the household in full, the total amount of taxes paid by the household is fully compensated by the subsidies that return to the household in the form of firm profits. To see this, define Π as net profits (which does not include the subsidies), and gross profits as net profits plus the output and labor subsidies:

$$\Pi^{gross} = \underbrace{\int_{g} \sum_{j \in g} \left(P_{jg} Y_{jg} - P_k K_j - P_{m,j} M_j \right) dg - \int_{s} \sum_{j \in s} \Phi_{js} H_{js} ds}_{\equiv \Pi} + \int_{g} \sum_{j \in g} T_y(Y_{jg}) dg + \int_{s} \sum_{j \in s} T_h(H_{js}) ds$$

Then, the household's budget constraint under the subsidies is:

$$C = \Phi H + \Pi^{gross} - T_y^{total} - T_h^{total}$$
$$= \Phi H + \Pi$$

where $T_y^{total} = \int_g \sum_{j \in g} T_y(Y_{jg}) dg$ and $T_h^{total} = \int_s \sum_{j \in s} T_h(H_{js}) ds$ are the lump sum taxes on the household. The subsidies therefore only affect the household budget constraint through firms' input demand conditions. These subsidies are meant as a tool to assess the aggregate welfare effects of firm market power. In a richer setting that considers household heterogeneity, the implementation of such taxes and subsidies can have direct distributional implications that my model does not address.

C.6 Calibrating markup and markdown-related parameters

C.6.1 Calibration details

Section 6.3 briefly describes the calibration of the preference parameters (θ , σ , ν , η) underlying price markups and monopsony markdowns based on the model equations:

$$\begin{array}{ll} \text{Markups:} & \frac{\mu_{jg} - 1}{\mu_{jg}} = \frac{1}{\sigma} + \left(\frac{1}{\theta} - \frac{1}{\sigma}\right) \frac{P_{jg}Y_{jg}}{\sum_{j'}^{n_g} P_{j'g}Y_{j'g}}\\ \text{Markdowns:} & \frac{\lambda_{js}}{1 - \lambda_{js}} = \eta + (\nu - \eta) \frac{\Phi_{js}H_{js}}{\sum_{j'}^{n_s} \Phi_{j's}H_{j's}} \end{array}$$

In this section, I describe the implementation of these equations in more detail and assess how well the variable markups implied by these equations fit the relationship between my empirical markup estimates and measured market shares.

Product demand parameters. To calibrate the within-sector elasticity of substitution between varieties (σ), I first measure product market shares as the share of sales within 5-digit sectors. I then run a pooled OLS regression of the inverse price elasticity of demand ($\hat{\rho}_{jg}^{-1}$)—computed using the estimated markups—on market shares, controlling for market×year fixed effects. This gives an estimate for the passthrough of market shares to markups ($\hat{1}_{\theta} - \hat{1}_{\sigma}$). I compute an economy-wide σ as the median of the inverse of $\hat{\rho}_{jg}^{-1} - (\hat{1}_{\theta} - \hat{1}_{\sigma}) \frac{P_{jg}Y_{jg}}{\sum_{j'}^{n_g} P_{j'g}Y_{j'g}}$. To calibrate the between-sector elasticity of substitution (θ), I compute the inverse of ($\hat{1}_{\theta} - \hat{1}_{\sigma}$) + $\hat{1}_{\sigma}$ using the calibrated σ . Appendix C.6.2 discusses the measurement of market shares and presents values of σ by 2-digit sectors.

Labor supply parameters. Analogous to the calibration of the product demand parameters, I run a pooled OLS regression of the firm-specific labor supply elasticity $(\hat{\xi}_{js}^{-1})$ —computed using the estimated monopsony markdowns—on labor market shares, controlling for market×year fixed effects. I measure labor market shares as the share of wage bills within 5-digit sectors×commuting-zone pairs. This gives a market share passthrough estimate $(\nu - \eta)$. Next, I compute an economy wide η as the median of $\hat{\xi}_{js}^{-1} - (\nu - \eta) \frac{\Phi_{js}H_{js}}{\sum_{j'}^{n_s} \Phi_{j's}H_{j's}}$. To obtain ν , I measure $(\nu - \eta) + \eta$ using the calibrated value of η . In Appendix C.6.2, I find similar parameter values for η and ν when I widen the definition of a local labor market to a 2-digit sector×commuting-zone.

Model fit. The nested-CES product demand (labor supply) structure imposes a specific

relationship between markups (markdowns) and market shares. The first panel of Figure 19 shows how the estimated markups and the calibrated nested-CES markups vary by sales market share quantiles. The orange line shows that the estimated markups increase particularly steeply with market shares among the top two quantiles. The calibrated nested-CES markups are generally able to capture this pattern, particularly the markups of firms in the top quantile of market shares.



Figure 19: Nested-CES markups and monopsony markdowns.

Notes: This figure shows how estimated markups (markdowns) and nested-CES markups (markdowns) vary by sales (wage-bill) market share. Markup estimates, markdown estimates, and market share measures are residualized by the market×year fixed effects. Firms in the top 5% of the market share distribution are in bin 20. The numbers in square brackets are the median market share within each bin.

The second panel of Figure 19 shows how the estimated monopsony markdowns and the calibrated nested-CES markdowns vary by wage-bill market share quantiles. Both the estimated markdowns and the calibrated markdowns are relatively flat across the market share quantiles.

C.6.2 Alternative measures of product and labor market shares

The nested-CES model presented in Section 6 imposes strong restrictions on markups and markdowns. First, markups and markdowns are a function of sales market shares and wage-bill market shares. Second, their relationship with market shares follow a specific functional form. In this section, I discuss the measurement issues regarding market shares and how they may affect the inferred preference elasticities and welfare calculations. My discussion below maintains the following assumptions: (i) product demand and labor supply preferences are nested-CES, and (ii) product and labor market structures are oligopolistic and oligopsonistic.

Market share measurement. The implementation of the markup equation implied by the nested-CES model requires measuring market shares for a given definition of product markets. Defining a product market is not straightforward, as researchers do not generally observe the direct competitors of a firm. I define a product market as a 5-digit manufacturing sector, as is standard in the literature (De Ridder et al., 2021; Amiti and Heise, 2024). There are 211 5-digit sectors. Under this definition, I use firm-level total revenues to compute market shares in my baseline implementation in Section 6.3.

There are two main challenges to measuring market shares in a world with international trade. First, if firms export across national borders, their exports do not compete directly with the output of domestic firms within the same market. Not accounting for exports may therefore overstate the market share of large firms. To address this concern, in Table 24, I report the values of the preference parameters (θ , σ) that I obtain when I measure market shares using domestic revenues. This is possible because the FARE firm balance sheet data breaks down firm-level total sales into domestic sales and exports, although the data does not contain any additional information about exports (e.g. destination). Overall, I find similar parameter values for (θ , σ) compared to my baseline calibration.

Second, market share measures that do not account for import competition will overstate the market share of domestic firms.³⁵ I now show that, in the presence of import competition, my calibration approach may understate degree of passthrough

³⁵See, for example, Amiti and Heise (2024), who show that accounting for foreign competition flattens the trend of rising market concentration in the US economy.

of market shares to markups $(\frac{1}{\theta} - \frac{1}{\sigma})$, and overstate the between-market elasticity of substitution (θ). To see this, let the true market share of a firm j in sector g be $s_{jgt}^{y} = \frac{P_{jgt}Y_{jgt}}{\sum_{j'}^{n_g} P_{j'gt}Y_{j'gt} + P_{Fgt}Y_{Fgt}}$, where $P_{Fgt}Y_{Fgt}$ are competing imports from foreign firms, and the measured market share be $\hat{s}_{jgt}^{y} = \frac{P_{jgt}Y_{jgt}}{\sum_{j'}^{n_g} P_{j'gt}Y_{j'gt}}$. The relationship between the true market share and the measured market share is $s_{jgt}^{y} = \hat{\gamma}_{gt}\hat{s}_{jgt}^{y}$, where $\hat{\gamma}_{gt} \equiv \frac{\sum_{j'}^{n_g} P_{j'gt}Y_{j'gt}}{\sum_{j'}^{n_g} P_{j'gt}Y_{j'gt} + P_{Fgt}Y_{Fgt}} \in [0, 1]$. The relationship between markups and market shares can then be written as:

$$\frac{\hat{\mu}_{jg}-1}{\hat{\mu}_{jg}} = \frac{1}{\sigma} + \left(\frac{1}{\theta} - \frac{1}{\sigma}\right)\hat{\gamma}_{gt}\hat{s}_{jgt}^{y}$$

Therefore, the estimated market share passthrough becomes $(\frac{1}{\theta} - \frac{1}{\sigma})\hat{\gamma}_{gt}$, which is less than $(\frac{1}{\theta} - \frac{1}{\sigma})$. Understating the market share passthrough does not necessarily affect the inferred σ , because σ is measured using $\frac{1}{\sigma} = \frac{\hat{\mu}_{jg}-1}{\hat{\mu}_{jg}} - (\frac{1}{\theta} - \frac{1}{\sigma})\hat{\gamma}_{gt}\hat{s}_{jgt}^y$. However, it would lead to an upward-bias in the inferred θ , because this parameter is measured as the inverse of $(\frac{1}{\theta} - \frac{1}{\sigma})\hat{\gamma}_{gt} + \frac{1}{\sigma}$.

How might the biases induced by unobserved import competition affect the welfare calculations in Section 7.4? Section 6.2 shows that the dispersion of markups reduce welfare by misallocating factor inputs across firms. In the nested-CES model, the lower bound for markups is $\mu_{min} = \frac{\sigma}{\sigma-1}$ and the upper bound is $\mu_{max} = \frac{\theta}{\theta-1}$. The responsiveness coefficient $\frac{1}{\theta} - \frac{1}{\sigma}$ governs whether a firm with a given market share is closer to μ_{min} or μ_{max} . Therefore, the downward bias in the responsiveness of markups to market shares $(\frac{1}{\theta} - \frac{1}{\sigma})$ and the resulting upward bias in the between-market elasticity of substitution (θ) reduces the possible extent of markup dispersion. In turn, this will understate the welfare cost of markups operating through the misallocation channel. In this case, one can read the welfare costs of markups in Section 7.4 as a lower bound.

On the labor market side, the measurement of wage-bill market shares for local labor markets are less susceptible to issues concerning international trade. In my baseline calibration, I define a local labor market as a 5-digit sector and commuting zone pair. There are 211 5-digit sectors and 507 commuting zones in my sample, leading to 16,651 local labor markets. This approach follows Berger et al. (2022), who define local labor markets by 3-digit NAICS and commuting zone pairs. In Table 24, I also present the calibrated labor market preference parameters (ν , η) when I define local labor markets at the more aggregated 2-digit sector×commuting-zone level.

Within-market elasticities of substitution	σ		η	
Manufacturing sectors	(1)	(2)	(3)	(4)
Textile	3.45	3.45	0.00	0.00
Apparel	-3.11	-3.12	0.00	0.00
Leather	-7.19	-7.11	1.19	1.17
Wood products (excluding furniture)	6.82	6.74	0.68	0.72
Paper and publishing	4.44	4.41	0.82	0.87
Recorded media	2.43	2.43	1.60	1.50
Chemicals	41.24	7.88	0.77	0.66
Pharmaceutical	26.61	42.01	-0.06	0.16
Rubber & plastics	4.98	4.98	1.11	0.96
Non-metallic minerals	4.00	3.94	0.44	0.66
Basic metals	5.23	4.80	0.99	0.86
Fabricated metals (excluding machinery)	3.54	3.54	1.16	0.94
Computers, electronic, & optical	7.33	5.10	1.07	0.94
Electrical equipment	55.15	5.54	0.84	0.78
Machinery & equipment	13.34	12.42	0.62	0.55
Motor vehicles	-747.12	-13.95	0.84	0.82
Other transport equipment	8.69	8.68	0.02	0.16
Furniture	3.35	3.36	1.21	1.21
Other manufacturing	3.36	3.34	0.90	0.95
Repair & installation of machinery	4.29	4.42	0.97	0.84
Economy-wide	5.16	4.57	0.97	0.85
Between-market elasticities of substitution	θ		ν	
	1.21	1.23	0.67	0.40

Table 24: Elasticities of substitution between goods (σ) and between jobs (η) by 2-digit French manufacturing sectors.

This table reports calibrated values of the elasticities of substitution between goods (σ) and between jobs (η) for 2-digit sectors (2009-2016). The 2-digit sector-specific parameters are obtained by implementing the method described in Section 6.3 for each 2-digit sector. My baseline measures of (σ , η) are reported columns (1) and (3) in the row 'Economy-wide'. In columns (1) and (2), a product market is defined as a 5-digit sector. Market shares are measured using firm-level total sales in column (1), and as domestic sales in column (2). In column (3), a local labor market is defined as a 5-digit sector×commuting-zone pair. In column (4) a local labor market is a 2-digit sector×commuting-zone pair. Labor market shares are measured as a firm's wage-bill share in columns (3) and (4).

C.7 Measuring quality and amenities under variable elasticity of substitution

In Section 6.3, I measure firm heterogeneity in product quality and non-wage amenities under the assumption of CES preferences. An important concern is whether the measured degree of heterogeneity depends heavily on the assumed preference structure. In this section, I measure these sources of heterogeneity using other well-known preference structures that nest the CES as a special case. Specifically, I compare the CES-implied measures with two, more general, classes of variable-elasticity-of-substitution (VES) product demand and labor supply systems: (i) Pollak (1971) additively separable preferences and (ii) Kimball (1995) preferences.³⁶ My main finding is that the implied firm heterogeneity under these two alternative preference structures are closely correlated with my baseline measures based on the CES structure.

In what follows, I maintain the assumption that the product demand aggregator over product markets, and the labor supply aggregator over labor markets, are CES. That is, at the aggregate level, the model structure is as in Section 6. However, at the product and labor market-level, the aggregator over firms are now VES instead of CES. The VES preference structure implies that markups and markdowns are firm-specific and depends on firm size, even if market structures are monopolistic and monopsonistic.³⁷ As such, in this section I assume that firms behave monopolistically in product markets and monopsonistically in labor markets.

C.7.1 Pollak's additive preferences

Let the aggregator over firms (product varieties) in market *g* be:

$$Y_g = \sum_{j \in g} \tilde{D}_{jg}^{\frac{1}{\sigma}} \left(\frac{\sigma}{\sigma-1}\right) \left(Y_{jg} + \delta^y\right)^{\frac{\sigma-1}{\sigma}}$$

where $\sum_{j \in g} \tilde{D}_{jg} = 1$. The parameter $\delta^y \ge 0$ governs the extent of the departure from CES, with $\delta^y = 0$ giving the CES demand system. Solving for the expression for market shares gives:

³⁶These preferences have been used extensively in research in international trade and macroeconomics. See, for example, Morlacco and Arkolakis (2017), Arkolakis, Costinot, Donaldson, and Rodrigues-Clare (2019), and Edmond et al. (2023).

³⁷See Morlacco and Arkolakis (2017) for a detailed exposition of these demand systems.

$$\frac{P_{jg}Y_{jg}}{\sum_{j'\in g}P_{j'g}Y_{j'g}} = \frac{\tilde{D}_{jg}^{\frac{1}{\sigma}}(Y_{jg} + \delta^{y})^{-\frac{1}{\sigma}}Y_{jg}}{\sum_{j'g}\tilde{D}_{j'g}^{\frac{1}{\sigma}}(Y_{j'g} + \delta^{y})^{-\frac{1}{\sigma}}Y_{j'g}}$$
(29)

and the price elasticity of demand is:

$$\rho_{jg}^{Pollak} = \sigma \left(1 + \frac{\delta^y}{Y_{jg}} \right) \tag{30}$$

As $\delta^y \to 0$, the price elasticity of demand and markups converge to the CES case: $\rho_{jg}^{Pollak} \to \sigma$ and $\mu_{jg}^{Pollak} = \frac{\rho_{jg}^{Pollak}}{\rho_{jg}^{Pollak-1}} \to \frac{\sigma}{\sigma-1}$. When $\delta^y > 0$, the price elasticity of demand decreases with firm size, while markups increase with firm size.

One feature of Pollak preferences is that price elasticities of demand (ρ_{jg}^{Pollak}) are infinitely large, and so markups ($\mu_{jg} - 1$) are approximately zero, for small firms (see equation (30)).³⁸ That is, small firms are approximately pricetakers. This implication appears to be inconsistent with my baseline markup estimates in Table 1, as most firms charge markups significantly above zero. However, the nested-CES product demand system in Section 6 avoids this inconsistency, as even small firms can charge non-negligible markups in that set up.

Analogously, let the labor supply aggregator over firms in labor market *s* be:

$$H_{s} = \sum_{j \in s} \tilde{A}_{js}^{-\frac{1}{\eta}} \left(\frac{\eta}{\eta-1}\right) \left(H_{js} + \delta^{h}\right)^{\frac{1+\eta}{\eta}}$$

where $\sum_{j \in s} \tilde{A}_{js} = 1$. The parameter $\delta^h \ge 0$ governs the extent of the departure from CES. Solving for the labor market shares gives:

$$\frac{\Phi_{js}H_{js}}{\sum_{j'\in s}\Phi_{j's}H_{j's}} = \frac{\tilde{A}_{js}^{-\frac{1}{\eta}}(H_{js}+\delta^h)^{\frac{1}{\eta}}H_{js}}{\sum_{j's}\tilde{A}_{j's}^{-\frac{1}{\eta}}(H_{j's}+\delta^h)^{\frac{1}{\eta}}H_{j's}}$$
(31)

and the labor supply elasticity:

$$\xi_{js}^{Pollak} = \eta \left(1 + \frac{\delta^h}{H_{js}} \right) \tag{32}$$

When $\delta^h > 0$, larger firms face lower labor supply elasticities and are able to mark down wages below the marginal revenue product of labor further. Just as for markups,

³⁸Other popular VES preferences also share this feature, for example, the Kimball demand system.

the Pollak labor supply system implies that small firms face infinitely large labor supply elasticities, and therefore, pay wages that are approximately the full marginal revenue product of labor ($\lambda_{js}^{Pollak} = 1$). This implication is also inconsistent with my baseline markdown estimates in Table 3, which shows that most firms would have non-negligible monopsony power in the absence of worker bargaining power. By contrast, the nested-CES labor supply system in Section 6 avoids this inconsistency.

Calibration and comparison. To compare my quality and amenity measures based on the CES structure to those obtained under the Pollak VES structure, I must first calibrate the parameters $(\delta^y, \delta^h, \sigma, \eta)$. To do so, I use moments from my baseline markup and markdown estimates, $\hat{\mu}_{jg}$ and $\hat{\lambda}_{js}$. Symbols with a hat are estimates. I maintain the same definition of a product market (5-digit sector) and labor market (5-digit sector×commuting-zone) as in Section 6. Because the Pollak product demand and labor supply systems nest the CES system I use in Section 6, I maintain the same calibrated values of σ and η as in my baseline calibration.

To calibrate the extent of the departure from CES in the Pollak product demand system, δ^{y} , I use equation (30). I write this equation in relative terms, comparing a firm *j* to the median firm in the markup distribution:

$$\frac{\hat{\rho}_{jg}^{Pollak}}{\hat{\rho}_{\mu,pct50}^{Pollak}} = \left(1 + \frac{\delta^{y}}{Y_{jg}}\right) \left/ \left(1 + \frac{\delta^{y}}{Y_{\mu,pct50}}\right)$$
(33)

where $\hat{\rho}_{jg}^{Pollak} \equiv \frac{\hat{\mu}_{jg}}{\hat{\mu}_{jg}-1}$. The parameter δ^y is then calibrated to match the covariance between firm size Y_{jg} and ratio of price demand elasticities $\frac{\hat{\rho}_{jg}^{Pollak}}{\hat{\rho}_{\mu,pct50}^{Pollak}}$. This moment is informative about the degree of departure from CES in the Pollak demand system, because CES (in combination with monopolistic competition) implies that price demand elasticities should be independent of firm size.

Similarly, I use equation (32) to calibrate the extent of the departure from CES in the labor supply system. I express this equation in relative terms, comparing a firm j to the median firm in the markdown distribution:

$$\frac{\hat{\xi}_{js}^{Pollak}}{\hat{\xi}_{\lambda,pct50}^{Pollak}} = \left(1 + \frac{\delta^h}{H_{js}}\right) \left/ \left(1 + \frac{\delta^h}{H_{\lambda,pct50}}\right)$$
(34)

where $\hat{\zeta}_{js}^{Pollak} \equiv \frac{\hat{\lambda}_{js}}{1-\hat{\lambda}_{js}}$. I then calibrate the parameter δ^h to match the covariance between firm size H_{js} and ratio of labor supply elasticities $\frac{\hat{\zeta}_{js}^{Pollak}}{\hat{\zeta}_{jn}^{Pollak}}$.



Figure 20: Measured product quality and non-wage amenity under nested-CES and Pollak's additively separable (VES) preferences in 2016. Note: Each data point is a firm.

Finally, given the calibrated parameters (δ^y , δ^h , σ , η), I measure firm heterogeneity in product quality using equation (29), and heterogeneity in non-wage amenities using equation (31). Figure 20 shows how the CES-based and VES-based (Pollak) measures of firm heterogeneity in quality and amenities compare. The left-hand-side panel of the figure shows that CES-based and Pollak-based quality measures correlate strongly with each other—they are centered around the 45-degree line. Similarly, CES-based and Pollak-based amenity measures correlate strongly with each other. The larger CES-Pollak disparity for quality measures compared to amenity measures arises because my markup estimates increase more sharply with firm size, while markdown estimates decrease at a slower rate with firm size (see Figure 12).

C.7.2 Kimball preferences

Let the aggregator over firms (product varieties) in market *g* be:

$$1 = \sum_{j \in g} \tilde{D}_{jg} Y\left(\frac{Y_{jg}}{\tilde{D}_{jg}Y_g}\right)$$

where $\sum_{j \in g} \tilde{D}_{jg} = 1$. The function Y(.) is homothetic, satisfies Y(1) = 1, Y' > 0, Y'' < 0, and nests CES preferences as a special case. Solving for the expression for market shares gives:

$$\frac{P_{jg}Y_{jg}}{\sum_{j'\in g}P_{j'g}Y_{j'g}} = Y'\left(\frac{Y_{jg}}{\tilde{D}_{jg}Y_g}\right)\frac{Y_{jg}}{Y_g}\mathcal{D}_g^{-1}$$
(35)

where $\mathcal{D}_g \equiv \sum_{j \in g} \Upsilon'(\frac{Y_{jg}}{\tilde{D}_{jg}Y_g}) \frac{Y_{jg}}{Y_g}$ is an endogenous demand shifter that is determined in general equilibrium.

Operationalizing the Kimball preferences requires making an assumption on the functional form of Y(.). I follow the common approach of assuming that Y(.) is an upper incomplete gamma function (Klenow and Willis, 2016). That is:

$$Y(x) = 1 + (\sigma - 1) \exp\left(\frac{1}{\delta^{y}}\right) (\delta^{y})^{\frac{\sigma}{\delta^{y}} - 1} \left[\operatorname{Gamma}\left(\frac{\sigma}{\delta^{y}}, \frac{1}{\delta^{y}}\right) - \operatorname{Gamma}\left(\frac{\sigma}{\delta^{y}}, \frac{x^{\frac{\delta^{y}}{\sigma}}}{\delta^{y}}\right) \right]$$

where, with a slight abuse of notation, the parameter $\delta^y \ge 0$ now governs the degree of departure of Kimball preferences from the CES; $\delta^y \to 0$ takes the preferences to the CES limit. With this functional form, one can write the a firm's product market share as:

$$\frac{P_{jg}Y_{jg}}{\sum_{j'\in g}P_{j'g}Y_{j'g}} = \tilde{D}_{jg}\exp\left[\frac{1}{\delta^{y}} - \frac{1}{\delta^{y}}\left(\frac{Y_{jg}}{\tilde{D}_{jg}Y_{g}}\right)^{\frac{\delta^{y}}{\sigma}}\right]\frac{Y_{jg}}{\tilde{D}_{jg}Y_{g}}\mathcal{D}_{g}^{-1}$$
(36)

and the price elasticity of demand as:

$$\rho_{jg}^{Kimball} = \frac{\sigma}{1 - \delta^{y} \log\left(\frac{P_{jg}}{P_{g}} \mathcal{D}_{g}\right)}$$
(37)

As $\delta^y \to 0$, the price elasticity of demand and markups converge to the CES case: $\rho_{jg}^{Kimball} \to \sigma$ and $\mu_{jg}^{Kimball} = \frac{\rho_{jg}^{Kimball}}{\rho_{jg}^{Kimball}-1} \to \frac{\sigma}{\sigma-1}$. When $\delta^y > 0$, the price elasticity of demand decreases with firm size, while markups increase with firm size.

Similarly, one can model the aggregator over firms in each labor market *s* as Kimball: (H_{i})

$$1 = \sum_{j \in s} \tilde{A}_{js} \mathcal{L} \left(\frac{H_{js}}{\tilde{A}_{js} H_s} \right)$$

where $\sum_{j \in s} \tilde{A}_{js} = 1$. The function $\mathcal{L}(.)$ is homothetic, satisfies $\mathcal{L}(1) = 1$, $\mathcal{L}' > 0$, $\mathcal{L}'' > 0$, and nests CES preferences as a special case. However, operationalizing the implied equations for labor supply (elasticities) requires assuming a functional form

that satisfies these conditions; one that is symmetric to the upper incomplete Gamma function of Klenow and Willis (2016). This is a challenging task, and to the best of my knowledge, no analogous function for labor supply has yet been developed. As such, I focus on comparing the Kimball-implied measures product quality heterogeneity to those implied by CES below.

Calibration and comparison. As with the calibration of the Pollak preference parameters, I maintain the same calibrated value of σ as in Section 6.3, as well as the same definition of product and labor markets. To calibrate the extent of the departure from CES in the Kimball product demand system, δ^y , I rearrange the equation (37) for the price elasticity of demand as:

$$\frac{1}{\hat{\rho}_{jg}^{Kimball}} = \frac{1}{\sigma} - \frac{\delta^{y}}{\sigma} \log P_{jg} + \frac{\delta^{y}}{\sigma} \log \left(\frac{\mathcal{D}_{g}}{P_{g}}\right)$$
(38)

where $\hat{\rho}_{jg}^{Kimball} \equiv \frac{\hat{\mu}_{jg}}{\hat{\mu}_{jg}-1}$. To calibrate the parameter δ^y , I first estimate the ratio $\frac{\delta^y}{\sigma}$ by running an OLS regression of the inverse of the price elasticity of demand $\frac{1}{\hat{\rho}_{jg}^{Kimball}}$ on output prices log P_{jg} , treating the last term in the equation, $\frac{\delta^y}{\sigma} \log \left(\frac{\mathcal{D}_g}{P_g}\right)$, as a sector×year fixed effect. Then, given the estimate $\hat{\delta^y}/\hat{\sigma}$, I measure δ^y as $\sigma(\hat{\delta^y}/\hat{\sigma})$. Finally, given these calibrated preference parameters, I measure firm heterogeneity in product quality using equation (35). This requires simultaneously solving for the general equilibrium quantities P_g , Y_g , and \mathcal{D}_g .

Figure 21 shows that CES-based and Kimball-based quality measures are closely aligned with each other, as they are approximately centered around the 45-degree line. However, it is worth noting that the Kimball demand system implies a higher quality measure than the CES for high-quality firms, but a lower quality measure than the CES for low-quality firms. Thus, the Kimball preference system implies a larger dispersion in product quality than the CES.



Figure 21: Measured product quality under nested-CES and Kimball (VES) preferences in 2016.

D Appendix: Additional Tables and Figures



Figure 22: Firm wage premia by deciles of firm size and labor productivity.





Figure 24: Output prices, markups, and labor wedges by high-wage status.

Notes: This figure shows how output prices, price-cost markups, and labor wedges vary by deciles of firm wage premia relative to firms in the first decile (low-wage firms). These are unconditional correlations. Decile 10 represents high-wage firms. Confidence intervals at the 95% level are plotted.



Figure 25: Output prices, markups, and labor wedges by deciles of firm size.

Notes: This figure shows how output prices, price-cost markups, and labor wedges vary by deciles of firm size (measured as sales) relative to firms in the first decile. These are unconditional correlations. Decile 10 represents large firms. Confidence intervals at the 95% level are plotted.

Manufacturing sectors	к	λ
Apparel	0.27	0.21
Repair & installation of machinery	0.24	0.27
Leather	0.21	0.49
Other transport equipment	0.18	0.25
Motor vehicles	0.16	0.38
Machinery & equipment	0.16	0.29
Chemicals	0.16	0.28
Other manufacturing	0.15	0.49
Textile	0.13	0.49
Wood products (except furniture)	0.13	0.41
Pharmaceutical	0.12	0.12
Furniture	0.11	0.58
Fabricated metals (except machinery)	0.10	0.48
Recorded media	0.08	0.63
Computers, electronic, & optical	0.08	0.52
Electrical equipment	0.08	0.49
Non-metallic minerals	0.08	0.46
Rubber & plastics	0.07	0.52
Basic metals	0.05	0.49
Paper and publishing	0.04	0.53

Table 25: Estimates of worker bargaining power and monopsony markdowns by 2-digit French manufacturing sectors.

This table reports estimates of worker bargaining power and median monopsony markdowns for 2-digit sectors (2009-2016).



Figure 26: Measured firm heterogeneity by high-wage status.

Notes: This figure shows how measured TFPQ, product quality, and non-wage amenities vary by deciles of firm wage premia relative to firms in the first decile (low-wage firms), controlling for 5-digit sector \times year fixed effects. All variables are in logs. Decile 10 represents high-wage firms. Confidence intervals at the 95% level are plotted.



Figure 27: Monopsony markdowns over time (2009-2016).

Table 26: Competition and the direct component of shock passthrough.

Shocks	$\sigma = 1.1$	$\sigma = 5$	$\sigma = 10$	$\eta = 1$	$\eta = 5$	$\eta = 10$
Δ TFPQ	0.05	1.05	1.48	1.07	0.53	0.32
Δ product quality	0.50	0.26	0.16	0.26	0.13	0.08
Δ amenity	-0.50	-0.26	-0.16	-0.26	-0.13	-0.08

This table reports the size of the *direct* component of shock passthrough under different degrees substitutability of product varieties and jobs. The passthrough measures are with respect to a positive 1% shock. When varying σ in the first three columns, η is kept at the baseline calibrated value of 0.97. When varying η in the last three columns, σ is kept at the baseline calibrated value of 5.16.

E Appendix: Extension with Two Skill Types

One limitation of the current analysis is that the model does not feature skill heterogeneity. Allowing for heterogeneous skill may shed light on the extent of differences in bargaining power and monopsony power by skill. Moreover, if skill compositions differ across firms, then my baseline model may mismeasure labor wedges, subsequently affecting the bargaining power and markdown estimates. In this section, I extend the analysis in Section 5 to a setting with two skill types—high-skill and low-skill occupations. This extension allows the composition of skills to differ across firms. My findings suggest that my baseline estimates in Section 5, which do not allow for skill heterogeneity, masks interesting differences in bargaining power (higher for high-skill workers) and monopsony markdowns (lower for high-skill workers) between skills, but captures the firm-level average markdown and labor wedge well.

E.1 Model environment

Skill-specific labor supply. Let the subscript $o = \{0, 1\}$ denote high-skill (o = 1) and low-skill (o = 0) labor, and let Φ_{oj} be the firm-specific wage premium for skill o. A worker i with efficiency E_i and skill o obtains a wage $W_{ioj} = E_i \Phi_{oj}$. Efficiency units of labor of skill o in the firm is $H_{o,j} = \overline{E}_{o,j}L_{o,j}$, where $\overline{E}_{o,j}$ denotes average efficiency and $L_{o,j}$ denotes amount of labor. The upward-sloping labor supply curve facing each firm is $H_{o,j} = \mathcal{H}_o(\Phi_{o,j}, A_{o,j})$. The skill and firm-specific labor supply shifter $A_{o,j}$ represents the non-wage amenities at a firm j for skill o and is taken as given by the firm.

Product demand. As in the baseline model, the downward-sloping demand curve for firm *j*'s output is $Y_j = \mathcal{G}(P_j, D_j)$, where P_j is the price and D_j is a demand shifter (product quality), which is taken as given by the firm.

Production technology. Firms operate a general production function with diminishing marginal returns to each input $Y_j = \Omega_j F(K_j, M_j, H_{0j}, H_{1j})$, where the main change compared to my baseline model is that the production function now features two skill types. This production function allows skill compositions to differ across firms due to (i) skill-specific complementarities or substitutabilities with non-labor inputs, and (ii) non-homotheticities in skills.

As before, Ω_i is the Hicks-neutral productivity term, K_i are physical capital, and

 M_j are material inputs. I maintain the assumption that the markets for capital and material inputs are perfectly competitive with prices P_k and P_m .

Wage determination. Workers of each skill type bargain collectively with their employer *j*. I maintain the assumption that bargaining is efficient—workers and firms jointly decide on wages, prices, materials, and capital to maximize total profits, taking into account the product demand curve and labor supply curves.

Denote high-skill workers' bargaining power as κ_1 and low-skill workers' bargaining power as κ_0 . The employer's bargaining power is then $1 - \kappa_0 - \kappa_1$. As before, in the event that firms and their employees do not arrive at an agreement, I assume that firms do not produce and have an outside option of zero profits; workers do not supply labor and have an outside option of zero wages. The worker-firm bargaining problem is:

$$\max_{\Phi_{0j},\Phi_{1j},P_{j},M_{j},K_{j}} \left(\Phi_{1j}H_{1j}\right)^{\kappa_{1}} \left(\Phi_{0j}H_{0j}\right)^{\kappa_{0}} \left(\Pi_{j}\right)^{1-\kappa_{0}-\kappa_{1}}$$

subject to $H_{oj} = \mathcal{H}_o(\Phi_{oj}, A_{oj}), Y_j = \mathcal{G}(P_j, D_j)$, and $Y_j = \Omega_j F(K_j, M_j, H_{0j}, H_{1j})$. The firm's profit is $\Pi_j = P_j Y_j - \Phi_{0j} H_{0j} - \Phi_{1j} H_{1j} - P_m M_j - P_k K_j$.

Skill-specific labor wedges. Solving the worker-firm bargaining problem gives the following skill-specific labor wedge equation at firm *j*:

$$\Lambda_{o,j} = \kappa_o \left(1 - \frac{P_k K_j + P_m M_j + \Phi_{o',j} H_{o',j}}{P_j Y_j} \right) \frac{\mu_j}{\alpha_{o,j}} + (1 - \kappa_0 - \kappa_1) \lambda_{o,j}$$

$$= \kappa_o \left(1 - \frac{\alpha_{k,j} + \alpha_{m,j} + \alpha_{o',j} \Lambda_{o',j}}{\mu_j} \right) \frac{\mu_j}{\alpha_{o,j}} + (1 - \kappa_0 - \kappa_1) \lambda_{o,j} \text{, s.t. } o \neq o'$$
(39)

where the α 's are output elasticities, λ 's are monopsony markdowns, and Λ 's are labor wedges. This equation is analogous to the one in the baseline model (see equation (3)). As the bargaining power of workers with skill *o* approaches zero, their labor wedge converges to the monopsony markdown. If, instead, workers of type *o* have some bargaining power, then they are able to capture some of the profits generated by markups, as well as profits from the employer's monopsony power over workers of the other skill type *o*'.

E.2 Estimating firm market power and worker bargaining power

Defining high-skill and low-skill occupations. I use the one-digit occupation classifications in the DADS matched employer-employee data set to define low-skill occupations as blue-collar occupations (e.g. maintenance workers and welders) and administrative support occupations (e.g. clerical workers and secretaries); I define high-skill occupations as senior staff in top management positions (e.g. head of logistics or human resources), employees in supervisory roles (e.g. accounting and sales managers), and technical workers (e.g. IT and quality control technicians).

Estimating firm wage premia. Let the subscript $o = \{1, 0\}$ denote high-skill and low-skill labor. The estimation procedure is as described in Section 3.1. However, firm-group effects are now occupation-specific:

$$\ln W_{it} = \iota_i + \phi_{g(j(i,t)),t} + \operatorname{Occ}_{o(i,t)} \times \phi_{g(j(i,t)),t} + \chi'_{it}\beta + \nu_{it}$$

where *i* denotes the individual, *j* denotes the firm, g(j) denotes the group of firm *j* at time *t*, o(i, t) denotes worker *i*'s occupational group at time *t*, ι_i are worker fixed effects, $\phi_{g(j(i,t))}$ are firm-group fixed effects, and χ_{it} is a vector of time-varying worker characteristics.

Estimating skill-specific labor wedges. Under the assumption that material inputs are flexible and that firms are pricetakers in material markets, the labor wedges can be measured as:

$$\Lambda_{o,j} = \left(\frac{\Phi_{o,j}H_{o,j}}{P_m M_j}\right) \left(\frac{\alpha_{m,j}}{\alpha_{o,j}}\right)$$
(40)

The key challenge to identifying the skill-specific labor wedge is to separately estimate them from skill-specific output elasticities.

To estimate the skill-specific output elasticities, I estimate a production function of the form:

$$y_j = \bar{\alpha}_0 f(k_j, m_j, h_{0,j}) + \bar{\alpha}_1 h_{1,j} + \omega_j$$

where, as a reminder, variables lowercase letters are in logs. This functional form assumes that high-skill labor are complements to the bundle of inputs comprised of capital, materials, and low-skill labor. This assumption is motivated by a large literature in macroeconomics on capital-skill complementarities (Krusell, Ohanian, Rios-Rull, and Violante, 2000), and on the complementarity between imported inputs and skills in international trade (Verhoogen, 2008). This Cobb-Douglas functional form in high-skill labor and the bundle of other inputs also mitigates the dimensionality problem associated with adding more factor inputs into the production function; with a flex-ible functional form such as the translog, the number of parameters to be estimated increases exponentially with the number of factor inputs. I approximate f(.) with a translog functional form. This accommodates the possibility that low-skill labor are substitutes with capital and materials.

Estimating this production function gives the high-skill output elasticity $\bar{\alpha}_1$ and low-skill output elasticity $\bar{\alpha}_0 \frac{\partial f}{\partial h_{0,j}}$. With these in hand, I back out the skill-specific labor wedges using equation (40).

Estimating skill-specific bargaining power and monopsony markdowns. My approach to estimating skill-specific worker bargaining power, which is motivated by the need to avoid the unobserved endogenous monopsony wages, is as described in Section 3.3. The estimating equation is equation (39), restated below for convenience:

$$\Lambda_{o,j,t} = \kappa_o \underbrace{\left(1 - \frac{P_{k,t}K_{j,t} + P_{m,t}M_{j,t} + \Phi_{o',j,t}H_{o',j,t}}{P_{j,t}Y_{j,t}}\right) \frac{\mu_{j,t}}{\alpha_{o,j,t}}}_{\equiv \tilde{\mu}_{o,j}} + (1 - \kappa_0 - \kappa_1)\lambda(H_{o,j,t}, A_{o,j,t}), \text{ s.t.}$$

where $o \neq o'$ and $\tilde{\mu}_{o,j}$ are profits that can be appropriated by workers of skill o (as a share of revenues). As before, the main challenge to disentangling worker bargaining power (κ_o) from monopsony markdowns ($\lambda_{o,j,t}$) is the presence of unobserved amenities ($A_{o,j,t}$). To address the unobserved amenities, I apply the same insight from Section 3.3 that firm size in terms of employment and wage-bills jointly controls for unobserved amenities, $\lambda(H_{o,j,t}, A_{o,j,t}) = \lambda(H_{o,j,t}, \mathcal{H}_o^{-1}(H_{o,j,t}, \Phi_{o,j,t}H_{o,j,t}))$.³⁹ As before, I approximate the function $\lambda(.)$ with a fourth-order polynomial. Given the estimated bargaining power parameters, monopsony markdowns are measured as $\lambda_{o,j,t} = \frac{\Lambda_{o,j,t} - \hat{\kappa}_0 \tilde{\mu}_{o,j,t}}{1 - \hat{\kappa}_0 - \hat{\kappa}_1}$.

³⁹Note that, even when amenities vary across firms, they do not necessarily affect monopsony markdowns directly. This is the case when amenities are multiplicatively separable from wages in workers' labor supply functions. Under multiplicative separability, the control function for amenities is not needed. See the discussion in Section 3.3.

Main results. Table 27 reports the estimated firm wage premia, labor wedges, and monopsony markdowns by skill type. Comparing the first and second rows shows that high-skill workers receive higher wages, but have similar wage dispersion. Comparing the third and fourth rows shows that the average level of labor wedges are similar for high-skill and low-skill workers. However, there is much more dispersion in labor wedges among the high-skill. The labor wedges of the high-skill are higher than those of the low-skill at the 75th percentile, but lower than those of the low-skill at the 75th percentile, but lower than those of the low-skill at the 25th percentile. The fifth row shows the firm-level weighted average labor wedge (Λ_f) , where the weights are the employment share of a skill type within the firm. The level of the average firm-level labor wedges are higher than my baseline estimates.

Variable	Mean	Median	25^{th} Pct	75 th Pct	Var
ϕ_1	3.17	3.16	3.08	3.25	0.02
ϕ_0	2.62	2.61	2.54	2.70	0.02
Λ_1	0.70	0.54	0.30	0.91	0.41
Λ_0	0.66	0.67	0.54	0.84	0.33
Λ_f	0.75	0.69	0.54	0.91	0.21
$\lambda_1^{'}$	0.36	0.27	0.05	0.56	0.23
λ_0	0.42	0.47	0.30	0.63	0.15
λ_f	0.48	0.47	0.31	0.64	0.11

Table 27: Summary statistics for estimated firm wage premia, labor wedges, and monopsony markdowns by skill type.

This table reports the summary statistics for estimated firm wage premia, labor wedges, and wage markdowns by skill group. Monopsony markdowns are computed using the estimated bargaining power parameters in specifications (2) and (4) in Table 28 as the baseline.

Table 28 reports the estimates of worker bargaining power by skill type. Comparing the first two columns to the last two columns, I find that high-skill workers obtain a larger share of firm profits than low-skill workers. The sixth and seventh rows of Table 27 report the distribution of monopsony markdowns received by each skill group. The level of monopsony markdowns for high-skill workers are lower than that for low-skill workers, although the labor wedge for high-skill workers is higher. The last row of Table 27 shows that average firm-level markdown is similar to my baseline estimates in Table 3, which do not distinguish between skills.

Finally, I compare how labor wedges and monopsony markdowns vary across deciles of firm wage premia. Firm wage premia are computed as within-firm weighted

averages of the skill-specific firm wage premium at a given firm, with weights being the employment share of a given skill-group in the firm. Figure 28 presents the findings. Consistent with the finding that high-skill workers have a higher bargaining power, the labor wedges of high-skill workers are increasing in wage premia more steeply than for low-skill workers. As such, the firm-level (weighted) average labor wedge is also increasing with firm wage premia, consistent with my baseline estimates in Figure 1. Figure 28 also shows that high-wage firms tend to have slightly more monopsony power than low-wage firms—monopsony markdowns mildly decrease with firm wage premia. The opposite appears to be the case for high-skill workers, as monopsony markdowns rise with firm wage premia. The firm-level (weighted) average monopsony markdown has a flat profile across firm wage premia, in line my baseline estimates in Figure 2.

Taken together, these results suggest that my baseline estimates approximate a firm-level average labor wedge and average monopsony markdown.

Labor wedge	Hig	h-skill	Low-skill		
	(1)	(2)	(3)	(4)	
$-\widetilde{\mu}_{o,j}$	0.139	0.104	0.065	0.085	
	(0.000)	(0.001)	(0.001)	(0.001)	
Sector×year fixed effects	\checkmark	\checkmark	\checkmark	\checkmark	
Firm fixed effects		\checkmark		\checkmark	
Number of firms	85,349	82,089	85,349	82,089	

Table 28: Estimated workers' bargaining power.

This table reports the estimated workers' bargaining power by skill group. Sector fixed effects are at the 5-digit level. Block bootstrapped standard errors are reported in parentheses.



Figure 28: Skill-specific labor wedges and monopsony markdowns high-wage status.

Notes: This figure shows how skill-specific labor wedges and monopsony markdowns vary by deciles of firm wage premia relative to firms in the first decile, controlling for 5-digit sector×year fixed effects. Firm wage premia are constructed at the firm level by computing weighted average of skill-specific wage premia within firms, with the weights being the employment share of each skill-group within the firm. Decile 10 represents high-wage firms. Vertical bars are 95% confidence intervals.

Bibliography

- Abowd, J., F. Kramarz, and D. Margolis (1999). High Wage Workers and High Wage Firms. *Econometrica* 67(2), 251–333.
- Amiti, M. and S. Heise (2024). U.S. Market Concentration and Import Competition. *The Review of Economic Studies.*
- Amiti, M., O. Itskhoki, and J. Konings (2014). Importers, Exporters, and Exchange Rate Disconnect. *American Economic Review* 104(7), 1942–78.
- Arkolakis, C., A. Costinot, D. Donaldson, and A. Rodrigues-Clare (2019). The Elusive Pro-Competitive Effects of Trade. *The Review of Economic Studies 86*, 46–80.
- Atkeson, A. and A. Burstein (2008). Pricing-to-Market, Trade Costs, and International Relative Prices. *American Economic Review* 98(5), 1998–2031.
- Berger, D., K. Herkenhoff, and S. Mongey (2022). Labor Market Power. American Economic Review 112(4), 1147–1193.
- Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica* 64(4), 841–890.
- Bhuller, M., K. Moene, M. Mogstad, and O. Vestad (2022). Facts and Fantasies about Wage Setting and Collective Bargaining. *Journal of Economic Perspectives* 36(4), 29–52.
- Bilal, A. and H. Lhuillier (2025). Outsourcing, Inequality, and Aggregate Output. *Working paper*.
- Blundell, R. and S. Bond (1998). Initial Conditions and Moment Restrictions in Dynamic Panel Data Models. *Journal of Econometrics* 87, 115–143.
- Bond, S. and M. Söderbom (2005). Adjustment Costs and the Identification of Cobb Douglas Production Functions. *IFS working paper*.
- Bonhomme, S., T. Lamadon, and E. Manresa (2019). A Distributional Framework for Matched Employer-Employee Data. *Econometrica* 87(3), 699–738.
- Cahuc, P., F. Postel-Vinay, and J. Robin (2006). Wage Bargaining with On-the-Job Search: Theory and Evidence. *Econometrica* 74(2), 323–364.
- Card, D., A. Cardoso, J. Heining, and P. Kline (2018). Firms and Labor Market Inequality: Evidence and Some Theory. *Journal of Labor Economics* 36(1), 13–70.
- De Loecker, J. and P. Goldberg (2014). Firm Performance in a Global Market. *Annual Review of Economics 6*, 201–227.
- De Loecker, J., P. Goldberg, A. Khandelwal, and N. Pavcnik (2016). Prices, Markups,

and Trade Reform. *Econometrica* 84(2), 445–510.

- De Loecker, J. and P. Scott (2024). Markup Estimation using Production and Demand Data: An Application to the US Brewing Industry. *Working paper*.
- De Loecker, J. and C. Syverson (2021). An Industrial Organization Perspective on Productivity. *Handbook of Industrial Organization IV*.
- De Loecker, J. and F. Warzynski (2012). Markups and Firm-Level Export Status. *American Economic Review* 102(6), 2437–2471.
- De Ridder, M., B. Grassi, and G. Morzenti (2021). The Hitchhiker's Guide to Markup Estimation. *Working paper*.
- Di Addario, S., P. Kline, R. Saggio, and M. Solvsten (2020). It Ain't Where You're From, It's Where You're At: Hiring Origins, Firm Heterogeneity, and Wages. *Working paper*.
- Edmond, C., V. Midrigan, and D. Xu (2023). How Costly are Markups? *Journal of Political Economy* 131(7).
- Eeckhout, J. and P. Kircher (2011). Identifying Sorting In Theory. *Review of Economic Studies* 78(3), 872–906.
- Flynn, Z., A. Gandhi, and J. Traina (2019). Measuring Markups with Production Data. *Working paper*.
- Fougere, D., E. Gautier, and S. Roux (2016). Understanding Wage Floor Setting in Industry-Level Agreements: Evidence from France. *IZA discussion paper no.* 10290.
- Gandhi, A., S. Navarro, and D. Rivers (2020). On the Identification of Gross Output Production Functions. *Journal of Political Economy* 128(8), 921–947.
- Garicano, L., J. Van Reenen, and C. Lelarge (2016). Firm Size Distortions and the Productivity Distribution: Evidence from France. *American Economic Review* 106(11).
- Jäger, S., B. Schoefer, S. Young, and J. Zweimüller (2020). Wages and the Value of Nonemployment. *Quarterly Journal of Economics* 135(4), 1905–1963.
- Kimball, M. (1995). The Quantitative Analytics of the Basic Neomonetarist Model. *Journal of Money, Credit and Banking* (27), 1241–1277.
- Klenow, P. and J. Willis (2016). Real Rigidities and Nominal Price Changes. *Economica* 83, 443–472.
- Klette, T. and Z. Griliches (1996). The Inconsistency of Common Scale Estimators When Output Prices are Unobserved and Endogenous. *Journal of Applied Econometrics* 11(4), 343–361.

- Kline, P., N. Petkova, H. Williams, and O. Zidar (2019). Who Profits from Patents? Evidence from Innovative Firms. *Quarterly Journal of Economics* 134(3), 1343–1404.
- Kline, P., R. Saggio, and M. Solvsten (2020). Leave-Out Estimation of Variance Components. *Econometrica* 88(5), 1859–1898.
- Krueger, A. and L. Summers (1988). Efficiency Wages and the Inter-Industry Wage Structure. *Econometrica* 56(2), 259–293.
- Krusell, P., L. Ohanian, J.-V. Rios-Rull, and G. Violante (2000). Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis. *Econometrica* 68(5), 1029– 1053.
- Le Barbanchon, T., R. Rathelot, and A. Roulet (2019). Unemployment Insurance and Reservation Wages: Evidence from Administrative Data. *Journal of Public Economics* 171, 1–17.
- Morlacco, M. and C. Arkolakis (2017). Variable Demand Elasticity, Markups, and Pass-Through. *Working paper*.
- Mortensen, D. and T. Vishwanath (1991). Information Sources and Individual Wage Outcomes. Northwestern University, Department of Economics and Kellogg School of Management, Discussion Paper No. 948.
- Naouas, A. and F. Romans (2014). La négociation salariale d'entreprise de 2004 à 2010. *DARES report*.
- Pollak, R. (1971). Additive Utility Functions and Linear Engel Curves. *The Review of Economic Studies* 38, 401–414.
- Rubens, M. (2023). Market Structure, Oligopsony Power, and Productivity. *American Economic Review* 113(9), 2382–2410.
- Sorkin, I. (2018). Ranking Firms Using Revealed Preference. *Quarterly Journal of Economics* 133(3), 1331–1393.
- Stansbury, A. and L. Summers (2020). The Declining Worker Power Hypothesis. *Brookings Papers on Economic Activity*.
- Stole, L. and J. Zwiebel (1996). Intra-Firm Bargaining under Non-Binding Contracts. *Review of Economic Studies* 63(3), 375–410.
- Taschereau-Dumouchel, M. (2020). The Union Threat. *Review of Economic Studies* 87(6), 2859–2892.
- Verhoogen, E. (2008). Trade, Quality Upgrading, and Wage Inequality in the Mexican

Manufacturing Sector. *Quarterly Journal of Economics* 123(2), 489–530.

Yeh, C., C. Macaluso, and B. Hershbein (2022). Monopsony in the US Labor Market. *American Economic Review* 112(7), 2099–2138.