Multi-unit Collective Bargaining in the Presence of Labor Market Power

Sebastián Poblete Coddou[†] Gustavo González[‡] Nicolás Rivera[§]

[†]Northwestern University [‡]PUC-Chile [§]Central Bank of Chile

UPDATED FREQUENTLY, CLICK HERE FOR MOST RECENT VERSION

Abstract

Collective bargaining is one of the primary institutions through which workers can counteract employers' monopsony power, but its effectiveness depends on how bargaining is structured. This paper studies multi-unit bargaining—unions that negotiate with several firms rather than a single employer—and its effects on wages, employment, and welfare. We develop and estimate an equilibrium model of wage determination through bargaining in which firms retain the right to manage. We show that multi-unit unions have an incentive to internalize "employment diversion": raising wages at one firm shifts employment away from the other firms they represent, which dampens their incentive to push wages up. We estimate the model on matched employer-employee and tax records from Chilean manufacturing, where single- and multi-unit bargaining coexist. Multi-unit unions hold substantially more bargaining power than single-unit unions, and this advantage dominates the dampening effect of employment diversion: equalizing bargaining power across union types and removing diversion internalization lowers wages by 3.6% and employment by 10.5% for workers covered by multi-unit unions. Bargaining power gains alone raise wages by 3.9% and employment by 11.4%, while employment diversion alone lowers wages by 0.3% and employment by 0.9%. Because most firms in our sample are supply-constrained, internalizing diversion pushes wages further below the competitive level, so the case for centralized bargaining depends on the balance of employer and union market power.

JEL Codes: J51, J42, J31, J23

Keywords: collective bargaining, labor market power, monopsony, wage determination, unions

Acknowledgements: I want to thank my thesis committee—Rob Porter, Gastón Illanes, Vivek Bhattacharya, and Abigail Adams—for their invaluable guidance and support. I also want to thank Mar Reguant, Elisa Jácome, David Benson, Elena Prager, and seminar participants at Northwestern University and the Central Bank of Chile for their insightful and useful comments. All errors are my own.

1 Introduction

There is growing evidence that labor markets are not perfectly competitive, and that firms can exercise monopsony power across different contexts and industries, pushing wages and employment below the socially efficient levels.[1] One of the main institutions that allow workers to mitigate the effects of monopsony power is collective bargaining through a union. By coordinating workers, a union exercises market power of its own, which countervails the market power of the firm (Robinson, 1933; Galbraith, 1954). This countervailing power can also overshoot: a union with enough market power can drive wages above the competitive level, causing employment rationing. Understanding what drives union market power is therefore key to knowing *when* collective bargaining helps achieve a socially efficient outcome. One such driver is the degree of "multi-unit" coordination in bargaining—i.e., whether bargaining happens at the firm level or through unions that represent workers across firms in sectoral, regional or other associations. While the degree of coordination varies substantially across countries (Bhuller et al., 2022), there is little evidence on how it affects wages, employment, and welfare.

In this paper, we study how multi-unit bargaining affects workers' wages, employment, and welfare in the Chilean formal manufacturing sector. The structure of bargaining matters for the average worker's outcomes through two main forces. The first is employment diversion. A union that bargains with more than one firm internalizes that jobs are substitutes for workers: once the labor market clears again, pushing up wages at one firm can lower employment at the other firms the union represents. Internalizing this "employment diversion" leads the union to bargain for lower wages than it would if it bargained with a single firm. This, in order to avoid employment losses elsewhere in its bargaining group[2]. The strength of this force depends directly on how substitutable jobs are for workers. The second force is bargaining power: multi-unit unions may hold more bargaining power than single-unit unions, allowing them to better counteract firms' market power. Our goal is to measure how much each of these forces—employment diversion and differences in bargaining power—shapes workers' labor market outcomes.

To answer this question, we develop and estimate a structural model of wage setting with

[1] See Robinson (1933), and more recently Berger et al. (2022); Gottfries and Jarosch (2023); Setzler (2025); Kroft et al. (2025); Kline (2025).

[2] This is true for any objective function of the union that includes the level of employment in the represented firms.

bargaining and monopsony power. The model has three main components: a labor supply system that generates firm-level monopsony power, a labor demand system determined by the marginal revenue productivity of labor, and Nash-in-Nash bargaining between firms and unions over wages. Together, these components let us separate the union's relative bargaining power from the firm's market power for each firm-union pair in the data, and quantify how much employment diversion multi-unit unions internalize. We then use the estimated model to quantify the impact of multi-unit bargaining on equilibrium labor market outcomes. Our main counterfactual removes both features of multi-unit bargaining at once, which yields its overall effect on equilibrium wages and employment. We then decompose this effect in two further counterfactuals. First, we equate the bargaining power of multi-unit unions to that of single-unit unions, which isolates the effect of differences in bargaining power across union types. Second, we change the bargaining structure of multi-unit unions to single-unit while keeping their bargaining power fixed, which isolates the effect of internalizing employment diversion. Finally, we perform a heterogeneity analysis to study how the different regimes correlate with firm characteristics, and to analyze distributional effects.

We estimate the model on a comprehensive data set on the Chilean formal manufacturing sector, built from several administrative sources. We observe matched employer-employee data with wages and worker characteristics, which we join with firm-level data on revenues, materials expenditure and capital. We also observe the unions involved in collective agreements and the timing of those agreements, matched to their firms. Observing workers, their employers and their wages together with production fundamentals allows us to estimate both labor supply and labor demand, and to compute the counterfactuals above while accounting for equilibrium responses. These sources of variation also help us address endogeneity concerns arising from selection into different types of unions: the model and data let us account explicitly for differences across firms and unions in productivity (via production function estimation) and in supply responses (via labor supply estimation).

We model bargaining between firms and unions as bilateral Nash-in-Nash bargaining, in which the union bargains over the total wage bill of the workers in the firms it represents and the firm bargains over profits. Firms retain the right to manage: once the bargained wage is posted, each firm makes its hiring decision given that wage and the employment decisions of the other firms in the local labor market. As in Azkarate-Askasua and Zerecero (2025), the bargained wage can rise above

the competitive wage in equilibrium, in which case employment becomes demand-constrained and the firm rations employment. The Nash-in-Nash solution sets the bargained wage equal to the marginal revenue productivity of labor times a wedge that is weakly below one; this wedge is a weighted average of an endogenous oligopsony markdown and a union markup. The weights represent the union market power, which acts as a countervailing force to the firm's market power.

In our model, multi-unit bargaining partly dampens union market power, because employment diversion reduces the union's market power in equilibrium. The size of this reduction depends on the sum of the cross-elasticities of labor supply between the current firm and the other firms the union represents. Whether this dampening improves welfare depends on which side of the market holds more power. If the firm is demand-constrained, employment diversion leads the union to bargain for lower wages, bringing the outcome closer to the socially efficient one. If instead firm market power is high and the wage is below the competitive level, employment diversion leads the union to acquiesce and bargain for even lower wages than in the single-unit case, holding bargaining power constant.

The model's equilibrium hinges on endogenous labor supply and demand responses, together with the firm's own rationing decision. On the labor supply side, we estimate a nested logit model of labor supply with two main features. First, nesting workers' job choice by industry lets preferences for jobs be correlated within industry, so that substitution patterns across jobs are flexible rather than governed by a single market-wide elasticity. Second, because firms can limit how many of their applicants they actually hire, labor supply to a firm depends not only on its wage but also on its rationing rate, and supply responds endogenously as rationing rates change across firms. We recover these rationing rates jointly with the structural parameters of the model, so that labor supply and the firm's hiring decision remain mutually consistent in equilibrium. To achieve identification, we instrument the wage with the demographics of workers in other markets where the firm is posting jobs, and the nesting parameter with the number of competing firms in the same industry and market, which shifts a job's within-industry share without directly affecting the unobserved amenities workers value.

On the labor demand side, the firm maximizes profits taking into account the labor supply it faces. We estimate the marginal revenue productivity of labor with a production function approach, which gives us a labor demand curve for each firm. In the bargaining model, this demand curve

characterizes the marginal rents the firm captures in bargaining, as well as how those rents respond in equilibrium in the counterfactual exercises.

With labor supply and demand estimated, we recover the unions' relative bargaining power from the first-order conditions of the Nash-in-Nash problem, and the diversion effect from the estimated cross-elasticities of labor supply. We find that multi-unit unions have more bargaining power than single-unit unions, with a mean of 0.31 against 0.23, which generally translates into higher wages and employment for the workers they represent. In our main counterfactual, we shift the distribution of bargaining power of multi-unit pairs so that its mean matches that of single-unit pairs, and we remove the internalized diversion, effectively converting multi-unit unions into single-unit unions. Comparing the new equilibrium with the baseline, firms bargaining with multi-unit unions have, on average, 10.5% more employment and 3.6% higher wages.

The two decomposition counterfactuals separate this effect into bargaining power gains and diversion internalization. Equating the bargaining power of multi-unit unions to that of single-unit unions, while keeping diversion internalized, shows that the bargaining power gains of multi-unit unions translate into 3.9% higher wages and 11.4% higher employment. Removing the internalization of diversion for multi-unit pairs, while keeping their bargaining power fixed, shows that internalizing diversion lowers wages by 0.3% and employment by 0.9%, partially offsetting the bargaining power gains. Bargaining power gains are therefore quantitatively more important than the internalization of employment diversion.

In the heterogeneity analysis, we correlate bargaining power and multi-unit status with firm characteristics, and analyze distributional effects. Multi-unit unions are more likely to make firms demand-constrained, since they have more bargaining power and are sorted into firms with less elastic labor demand. We find no evidence that multi-unit unions are more likely to be in high-productivity firms, which suggests that the gains from multi-unit bargaining are not driven by systematic sorting on size or productivity.

Taken together, the results suggest that local labor market conditions are key to the optimal organization of unions. On average, bargaining power gains outweigh the diversion effect, so multi-unit bargaining raises both wages and employment, but this average hides heterogeneity. In firms with less elastic labor demand, where multi-unit unions tend to concentrate in our setting, misaligned incentives can cause rationing in equilibrium; and even where firms are supply-constrained, the

diversion effect works against a more efficient outcome. Evaluating multi-unit bargaining therefore requires accounting for labor market conditions and for market power on both sides of the market.

Related Literature. This paper relates to several strands of literature, from industrial organization concerns in labor markets to the labor economics literature on collective bargaining and monopsony power. First, we add to the growing empirical literature on the effects of monopsony or oligopsony power on labor markets (Berger et al., 2022; Gottfries and Jarosch, 2023; Setzler, 2025; Kroft et al., 2025; Kline, 2025), and the long literature on the impact of unions on labor market outcomes (Lewis, 1963; Freeman and Medoff, 1984; Card, 1996; DiNardo et al., 1996; Lemieux, 1998; DiNardo and Lee, 2004; Farber et al., 2021; Lagos, 2025; Felix, 2026; Dodini et al., 2026). The union literature estimates how much unions raise wages, using cross-sectional comparisons, longitudinal data on workers who switch into or out of union jobs, or quasi-experiments such as close union elections. We contribute to both literatures by modeling bargaining between unions and firms with market power, and by comparing how different types of union coordination affect labor market outcomes.

Second, we contribute to the related literature on industrial organization concerns in labor markets. There is growing evidence linking concentration to the exercise of oligopsony power by firms across settings and industries (Prager and Schmitt, 2021; Arnold, 2019; Delabastita and Rubens, 2025; Rubens, 2023; Sharma, 2025). As oligopsony power becomes a more prominent antitrust concern (Naidu et al., 2018), we shed light on the role of unions in countervailing employer market power, and on how different types of union coordination can be more or less effective in doing so. In this sense, regulation of union coordination complements antitrust policy in the labor market.

Third, our paper relates to the literature on equilibrium effects of bargaining in vertical relations. The Nash-in-Nash framework has become a workhorse to study bilateral bargaining (Horn and Wolinsky, 1988; Collard-Wexler et al., 2019), and has been used to study policy counterfactuals that involve equilibrium effects and bargaining externalities. Relevant examples include automobile auctions (Larsen, 2021), cable television (Crawford and Yurukoglu, 2012; Crawford et al., 2018), health markets (Grennan, 2013; Gowrisankaran et al., 2015; Ho and Lee, 2017, 2019; Cuesta et al., 2019), pension markets (Bhattacharya and Illanes, 2026), and financial markets (Robles-Garcia, 2020; Brancaccio and Kang, 2025). We apply this framework to the labor market, where multi-unit unions

internalize a bargaining externality, and quantify the effects of this externality on labor market outcomes.

Finally, there is a nascent literature that uses structural models to quantify the equilibrium effects of union bargaining in the presence of monopsony power. Azkarate-Askasua and Zerecero (2025) study aggregate effects of union bargaining in France, and Angerhofer et al. (2026) study the equilibrium effects of union bargaining and alternative policy counterfactuals in the US public teacher sector. We add to this literature by studying multi-unit bargaining in a setting where single- and multi-unit unions coexist, and by quantifying how differences in bargaining power and employment diversion affect labor market outcomes. More broadly, our work informs the discussion of union policy as a potentially powerful regulatory tool in markets with employer power.

2 Model

To describe and analyze the economic forces behind multi-unit bargaining, we develop a wage-setting model where unions and firms bargain over wages in a local labor market. This is the model that we take to the data in Section 4. Unions represent heterogeneous workers across firms, where unions can have objective functions other than the utility of the workers. Unions can also represent workers in more than one firm. Both firms and unions take into account the labor supply and demand conditions when bargaining over wages, which allows market power to be exercised on both the firm and the union side.

In a local labor market, workers can offer labor to a firm $j \in \mathcal{J}$, where $\mathcal{J}$ is the set of firms in the market. Wages are determined by bilateral bargaining between a union u and a firm j, where the union represents a set of jobs $\mathcal{B}_u \subseteq \mathcal{J}$. The timing of the labor market is as follows: First, firms and unions bargain over wages, with firms setting employment after the wage is bargained. Then, workers choose which firm to offer labor to, given the wages and employment levels posted by the firms.

In each labor market, firms generate revenues defined by $R_j(L_j)$ with $R_j'(L_j) > 0, R_j''(L_j) < 0$, where L_j is the number of workers employed by firm j. At a given wage W_j, the firm does not pick employment directly off the labor supply curve; instead, it chooses a *rationing rate* $\psi_j \in [0, 1]$, the

share of the labor offered at W_j that it actually hires, so that realized employment is $L_j = \psi_j L_j^S(W_j)$. The firm's profit is thus

$$\pi_j(\psi_j, W_j) = R_j\Big(\psi_j L_j^S(W_j)\Big) - W_j\,\psi_j L_j^S(W_j),$$

and the firm chooses $\psi_j \in [0,1]$ to maximize $\pi_j(\psi_j, W_j)$ given W_j. Let $L_j^D(W_j)$ denote the firm's unconstrained profit-maximizing employment level at wage W_j, defined by $R_j'\Big(L_j^D(W_j)\Big) = W_j$ – the usual downward-sloping labor demand curve. Following Azkarate-Askasua and Zerecero (2025), the first-order conditions of the firm's problem, together with the bound $\psi_j \leq 1$, characterize the optimal rationing rate for a given wage W_j:

$$\psi_j \leq 1, \qquad \psi_j L_j^S(W_j) \leq L_j^D(W_j), \qquad (\psi_j - 1)\Big(L_j^D(W_j) - \psi_j L_j^S(W_j)\Big) = 0. \tag{1}$$

Intuitively, if the wage is low enough that the firm would like to hire more than is offered ($L_j^D(W_j) \geq L_j^S(W_j)$), it hires everyone who applies: $\psi_j = 1$ and employment is *supply-constrained*. If instead the wage is high enough that the firm's desired employment falls short of what is offered ($L_j^D(W_j) < L_j^S(W_j)$), the firm rations down to exactly its desired level: $\psi_j = L_j^D(W_j)/L_j^S(W_j) < 1$ and employment is *demand-constrained*.

Workers maximize an indirect utility function that depends on wages and other characteristics of the job $U_j \equiv U(W_1, ..., W_J; j)$ and choose to offer labor to the one firm that maximizes utility. If there is rationing ($\psi_j < 1$), we assume that firms hire workers randomly. Thus, workers get hired with probability ψ_j. We assume workers obtain 0 utility if they are not hired, so under rationing they maximize expected utility, $\psi_j U_j$. Labor supply curves L_j^S will depend on both wages and rationing rates, and are such that $\partial L_j^S/\partial W_j > 0$ and $\partial L_j^S/\partial \psi_j > 0$. Upward-sloping labor supply curves can come from different sources, such as job differentiation and employer concentration. We consider both explicitly in the empirical implementation in Section 4. Cross-elasticities of labor supply are assumed to be negative, consistent with jobs being substitutes across firms for the worker: $\partial L_k^S/\partial W_j < 0$ for $k \neq j$. We are assuming that the worker will choose *one* "main" labor relationship or one full-time job.

Unions indexed by u represent workers in bilateral bargaining. Denote by $\mathbf{W}^u, \mathbf{L}^u$ the vectors

of wages and employment levels across the bargaining unit respectively. The union's objective function

$$F_u\left(\mathbf{W}^u, \mathbf{L}^u\right)$$

is assumed to be an increasing function of the employment and wages across the bargaining unit.

Under these simple assumptions on the union's objective function we can illustrate that in the multi-unit case there are forces with opposite signs when it comes to raising wages. To see this, first call $\epsilon^{a,b}$ the partial elasticity between a and b, and note that the total elasticity of the union's objective function with respect to wages on a single job j is:

$$\phi^u = \frac{dF_u}{dW_j}\frac{W_j}{F_u} = \underbrace{\epsilon^{F_u,W_j}}_{\text{Own-wage effect}} + \underbrace{\epsilon^{F_u,L_j} \times \epsilon^{L_j,W_j}}_{\text{Own-employment effect}} + \underbrace{\sum_{k\neq j, k\in\mathcal{B}_u} \epsilon^{F_u,L_k} \times \epsilon^{k,j}}_{\text{Employment diversion effect}} \tag{2}$$

where ϵ^{L_j,W_j} is the elasticity of employment at j with respect to its own wage. Which curve determines employment depends on the firm's regime in equation (1). If the firm is supply-constrained ($\psi_j = 1$), employment moves along the labor supply curve and $\epsilon^{L_j,W_j} = e_j > 0$, the labor supply elasticity. If it is demand-constrained ($\psi_j < 1$), employment moves along the labor demand curve and $\epsilon^{L_j,W_j} = \varepsilon_j^D < 0$, the labor demand elasticity. A higher wage therefore raises employment, and with it the union's objective, only on the supply-constrained side. On the demand-constrained side it lowers employment, and the union gains from a higher wage only if the own-wage effect outweighs that loss. The last term, $\sum_{k\neq j,k\in\mathcal{B}_u} \epsilon^{F_u,L_k} \times \epsilon^{k,j}$, where $\epsilon^{k,j}$ is the elasticity of employment at k with respect to the wage at j, represents the internalized employment diversion of the multi-unit case, and it is absent in the single-unit case. When firm j is supply-constrained, the term is negative: as the wage at j increases, job j becomes more attractive and employment at the union's other firms k falls, since jobs are substitutes, which reduces the union's objective. Everything else equal, a multi-unit union then settles for lower wages than a single-unit union with the exact same characteristics. When firm j is demand-constrained, the direction of diversion can reverse. A higher wage still makes job j more attractive to workers, but the firm now hires fewer of them, since it moves down its labor demand curve and rations more. If the second effect dominates, fewer workers are absorbed by firm j, more remain available to the union's other firms, and the diversion term becomes positive. A multi-unit union bargaining with a rationing firm can

therefore have *stronger* incentives to raise the wage than a single-unit union.

The union-firm pairs solve bilateral Nash bargains independently and simultaneously for each firm $j \in \mathcal{B}_u$. Both the firm and the union assume that the other bargaining pairs have set their wages optimally, thus bargaining is Nash-in-Nash. Therefore, the bargained wage W_j corresponds to the solution to the following maximization problem:

$$W_j = \underset{W_j}{\operatorname{argmax}}\ \pi_j(\psi_j, W_j)^{\left(1-\gamma_j^u\right)} \times \left[F_u\left(\mathbf{W}^u, \mathbf{L}^u\right)\right]^{\gamma_j^u}$$

$$\text{s.t.} \quad \psi_j \leq 1, \qquad \psi_j L_j^S(W_j) \leq L_j^D(W_j), \qquad \left(\psi_j - 1\right)\left(L_j^D(W_j) - \psi_j L_j^S(W_j)\right) = 0$$

where the bargain is over the wage W_j only – firms retain the right to manage, so the rationing rate ψ_j is set by the firm, unilaterally and after the wage is bargained, to satisfy the above constraint set. A market equilibrium is thus a set of wages $\{W_j\}_{j\in\mathcal{J}}$ and rationing rates $\{\psi_j\}_{j\in\mathcal{J}}$ such that (i) each ψ_j solves the firm's rationing problem (1) given W_j, and (ii) each W_j is the Nash-in-Nash bargained wage above for firms whose job is represented by a union, or the oligopsony wage $W_j = \arg\max_{W_j} \pi_j(\psi_j, W_j)$ for firms with no union present. The parameter γ_j^u represents the relative bargaining weight of the union in the bargaining pair (j, u). We now characterize the solution to the Nash-in-Nash problem as in Azkarate-Askasua and Zerecero (2025), adding the extension to the multi-unit bargaining case. Calling $\beta_j^L = \partial \ln R_j / \partial \ln L_j$ the revenue elasticity of labor, the first-order condition of this problem takes a different form in each regime of equation (1).

Supply-constrained regime ($\psi_j = 1$). The bargained wage is the marginal revenue product times a wedge,

$$W_j = \lambda_j R_j'(L_j), \quad \lambda_j = \omega_j^u \frac{1}{\beta_j^L} + (1 - \omega_j^u)\mu_j, \tag{3}$$

where ω_j^u is a bargaining weight that depends on the union's bargaining power γ_j^u and on the elasticity of the union's objective ϕ^u in equation (2), evaluated at $\epsilon^{L_j, W_j} = e_j$, and μ_j is the markdown term, which depends on the labor supply elasticity e_j:

$$\omega_j^u = \frac{\gamma_j^u \phi^u}{\gamma_j^u \phi^u + (1 - \gamma_j^u)(1 + e_j)} \in [0, 1], \quad \mu_j = \frac{e_j}{1 + e_j}.$$

The wedge λ_j is a weighted average of the markdown $\mu_j < 1$, which represents firm-side market power, and a union markup term $1/\beta_j^L > 1$, since $\beta_j^L < 1$. This regime applies as long as the wedge is below one, so that the wage does not exceed the marginal revenue product.

Demand-constrained regime ($\psi_j < 1$). Here the wage equals the marginal revenue product, $W_j = R_j'(L_j)$ with $L_j = L_j^D(W_j)$, so there is no wedge. The firm is on its labor demand curve, so by the envelope theorem a higher wage lowers its profit by exactly its employment, $\partial \pi_j / \partial W_j = -L_j$, and the labor supply elasticity drops out of the firm's side of the bargain. The first-order condition instead pins down *where* on the labor demand curve the wage lands. It has a solution only if the union still gains from a higher wage even though employment falls, which for a single-unit union requires labor demand to be inelastic at the bargained wage, $\varepsilon_j^D > -1$. Appendix D derives the first-order condition in each regime.

Single- and multi-unit unions have different incentives to raise wages in a particular firm, a result we can explore by looking closely at the bargaining weight. The bargaining weight ω_j^u is increasing in both the union's bargaining power γ_j^u and the elasticity of the union's objective function ϕ^u. As discussed above, the elasticity of the union's objective function ϕ^u is decreasing in the number of firms in the bargaining unit, everything else equal. Therefore, as the number of firms in the bargaining unit increases, the union's incentives to raise wages in a particular firm decrease, everything else equal, resulting in a smaller bargaining weight ω_j^u. This is a key result of the model: multi-unit unions have weaker incentives to raise wages in a particular firm than single-unit unions, everything else equal.

Figure 1 depicts these two opposing forces within a single local labor market for firm j. The figure shows the labor demand implicitly defined by the marginal revenue product ($W_j = R_j'(L_j^D)$); the labor supply curve $L_j^S(W_j)$; and $MC_j(L_j)$, the marginal cost of labor implied by the upward-sloping supply curve. Two benchmark outcomes are marked: the oligopsony wage (W_j^O, L_j^O), where $MC_j(L_j)$ crosses labor demand and $W_j^O = \mu_j R_j'(L_j^O)$ is the standard monopsony markdown of equation (3); and the competitive wage (W_j^C, L_j^C), where labor supply meets labor demand directly. As equation (3) makes clear, the set of possible bargained outcomes—the purple locus—is not confined to the labor supply curve. For wages between W_j^O and W_j^C, employment is labor-supply constrained ($\lambda_j < 1$, $\psi_j = 1$ in equation (1)) and the locus runs along $L_j^S(W_j)$. Once the

wage reaches the competitive level, it stays there for an intermediate range of bargaining power (Appendix D). Beyond that range, and provided labor demand is inelastic, further increases in bargaining strength push the locus up along the labor demand curve, with the employer setting $\psi_j = L_j^D(W_j)/L_j^S(W_j) < 1$ to ration employment down to $L_j = L_j^D(W_j)$, remaining on $R_j'(L_j)$.

Where exactly a multi-unit union settles on this locus is determined by equation (2). Also, how bargaining power and diversion internalization move the wage towards the competitive solution depends on whether the bargained outcome is labor-supply or labor-demand constrained. Panel (a) of Figure 1 depicts a bargained wage W_j^B on the supply-constrained segment of the locus. There, bargaining power pulls the wage up the locus, toward W_j^C, since a higher W_j raises L_j and, through it, the union objective F_u. The employment diversion effect pulls in the opposite direction, back towards W_j^O, since the same wage increase diverts employment away from another firm $k \in \mathcal{B}_u$ and thereby lowers the objective F_u through the resulting drop in L_k.

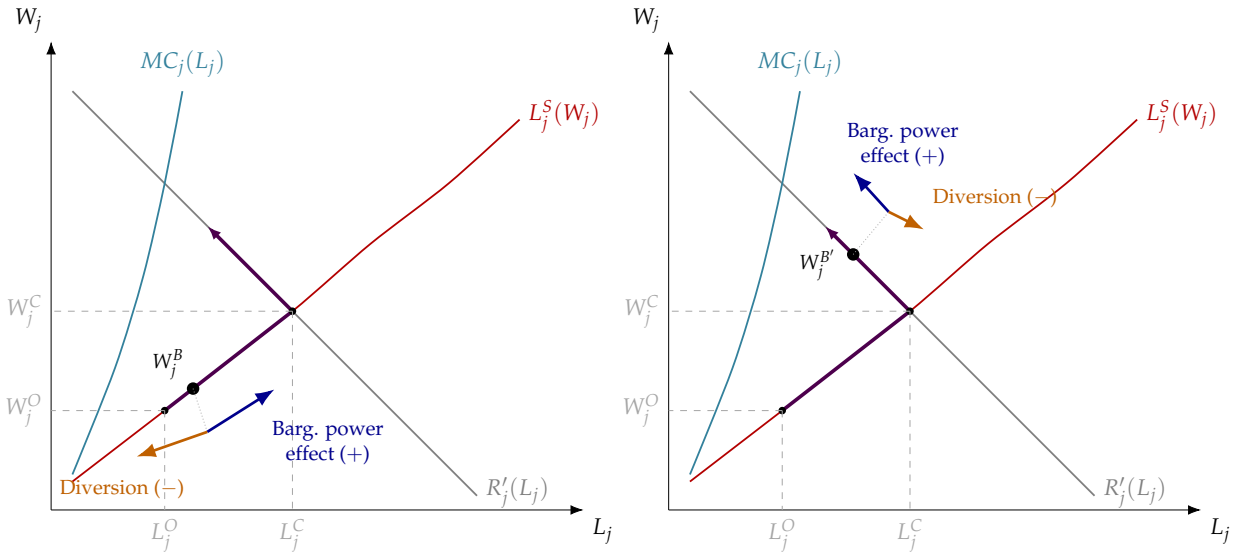
Panel (b) depicts a bargained wage $W_j^{B'}$ on the demand-constrained segment of the locus, beyond W_j^C. The two forces continue to pull in opposite directions: bargaining power pushes further up the locus, but beyond W_j^C that means pushing *away* from the competitive outcome, deeper into the rationing region. When the diversion term remains negative, the employment diversion effect still pushes back down the locus, and beyond W_j^C that now means pushing *toward* the competitive outcome. When it reverses sign, as discussed after equation (2), diversion instead reinforces bargaining power and pushes the wage further into rationing. How welfare changes as the bargained outcome moves along the locus is ambiguous: employment at j and at k move in different directions at different points on the locus, so the net effect depends on how labor is allocated across the whole market at each point, not only on the direction of the wage. At the same time, reallocation of labor due to bargaining will also depend on the market conditions, crucially, the slopes (elasticities) of the labor supply and demand curves.

3 Institutional Background and Data

3.1 Legal Framework and Data Sources

The legal regulations that govern collective bargaining in Chile inform some of the modeling choices we make to take the framework presented in Section 2 to the data. The process of bargaining is

Figure 1. Wage Bargaining in a Local Labor Market



(a) Supply-constrained **(b) Demand-constrained**

Notes: Both panels show the same labor market for firm j; the purple curve traces the locus of possible bargained outcomes as bargaining strength rises, from the oligopsony wage W_j^O (where $MC_j(L_j)$ crosses labor demand $R_j'(L_j)$) up along $L_j^S(W_j)$ to the competitive wage W_j^C, and then up along $R_j'(L_j)$ as the employer rations employment. Panel (a): at a supply-constrained wage W_j^B, the *bargaining power effect* (blue) pulls the wage up the locus while the *employment diversion effect* (orange) pulls it back down toward W_j^O—the two forces oppose each other. Panel (b): at a demand-constrained wage $W_j^{B'}$, the two forces still oppose each other, but their relationship to the competitive outcome reverses—the own-employment effect now pushes further away from W_j^C, deeper into rationing, while the employment diversion effect now pushes back toward W_j^C. A multi-unit union, which internalizes the diversion effect, therefore settles closer to W_j^O than a single-unit union in the supply-constrained region, but closer to W_j^C (i.e., less deep into rationing) in the demand-constrained region.

heavily directed and overseen by the law and the dispositions of the labor relations authority, the *Dirección del Trabajo* (DT). Workers and firms can have an informal bargaining process, which requires reporting the agreement to the DT and is governed by more general rules, or they can have a formal process, which is regulated in every step of negotiations and requires institutional oversight. The formal process gives protections to both parties in exchange for the oversight, in particular strike and job protection for workers and lockout permissions for firms. Both processes require workers to form a collective association or go through a union, and present an initial offer to the firm. Only workers in non-managerial positions can participate in collective agreements. The firm is also required to respond to the offer within 15 days, and then the parties have different rules for the following counteroffers depending on whether they are in the formal or informal process. Both processes closely resemble a sequential bargaining game, which supports the Nash-in-Nash framework we use in the model.

Another implicit assumption of the model supported by the Chilean legal framework is that the bargaining process is not coordinated on the firm side. In Chile, firms cannot explicitly coordinate their bargaining offers. In the multi-unit case, the union must bargain with each firm separately. This is relevant for our model, as it is consistent with the assumption that each bargaining event is independent in the sense that the outcome of other bargaining events is just assumed to be the Nash bargain solution.

The data used in the paper come from several administrative sources. Matched employer-employee data with wages and worker characteristics come from the Unemployment Insurance (UI) registry. The registry is essentially a census of the Chilean formal sector, as workers are required to participate in the UI and firms are required to report workers' salaries and characteristics every month to the registry. Data on firms' production fundamentals like revenues, expenditures on materials, and valuation of capital and investments come from tax records, in particular the value-added tax form F29, which firms are required to file monthly to register sales, investment, and value-added tax accounting. We complement these data with the Chilean Census and the National Employment Survey (a survey representative at the region level) to gather information on people outside of the labor market and on local labor markets, such as the population actively looking for jobs.

3.2 Descriptive Facts

Figure 2 shows the recent evolution of collective bargaining activity, split into agreements signed under the formal and informal processes described above, plus strikes started each year (strikes are only available under the formal process, which is the one that carries the legal right to strike). Strikes remain a small fraction of overall bargaining activity in a typical year, with the main exception being the spike around the 2019 social unrest. A crucial assumption of the model is that there is a very small probability of breakdown in the bargaining, since there have to be gains from trade from the process. The model would need to incorporate a substantial friction, like information asymmetries or coordination problems—which would be out of scope and worsen tractability—to explain a high probability of breakdown. The low frequency of strikes in the data reassures us that the model we are using is adequate for this setting.

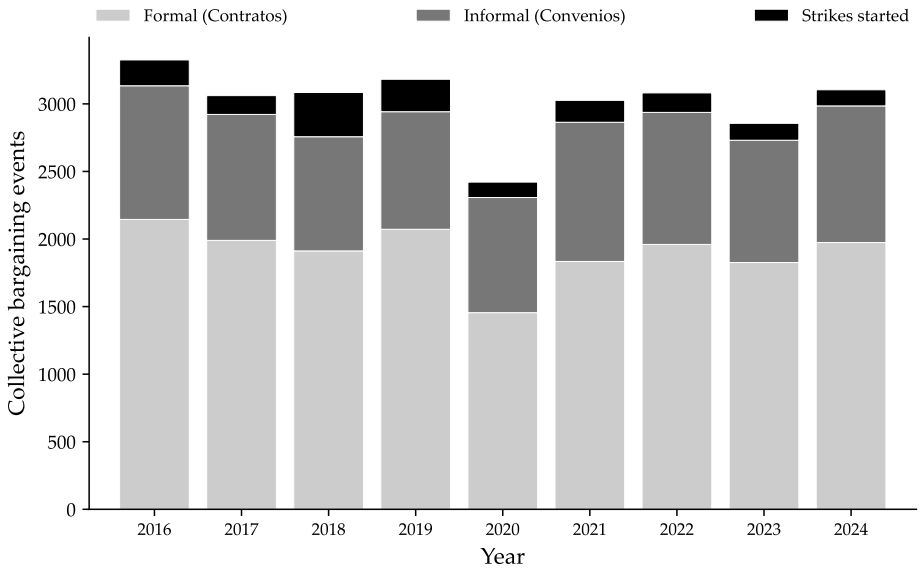


Figure 2. Collective bargaining activity in Chile, 2016–2024. Agreements signed each year under the formal (*Contratos Colectivos*) and informal (*Convenios Colectivos*) processes, with strikes started each year stacked on top. Source: Dirección del Trabajo.

The union coverage and membership in Chile in the studied period (2016-2025) is between 16% and 20%. This level is higher than in the US—around 6% in the same period (BLS, 2025)—and close to the average across OECD countries. Workers have the right to access the benefits of the collective agreement even if they are not union members. This mechanism is called *extensión de beneficios* (benefit extension) and only requires workers to pay union dues. Total coverage including workers

in benefit extension is between 19% and 20% of workers in the formal sector, which corresponds to between 1 and 1.3 million workers covered.

To estimate and compute counterfactuals using our labor supply model, we need to define *what* a local labor market is, and *who* is part of it. We define a local labor market to be the active population in a geographical commuting zone. To compute the local labor market sizes, we take the working-age population of each commuting zone and multiply it by the share of that population that is active—i.e., working or looking for a job—at the time. The proportion of active population in a market comes from the representative National Employment Survey (2026).

In terms of sample selection, we restrict the sample to manufacturing firms with at least 8 employees, which is the minimum size for a firm to be able to bargain collectively. Jobs at smaller firms are therefore part of the outside option. We also restrict the sample to firms that have positive revenue, expenditures in materials and capital, and positive employment throughout the period. We end up with a sample of 4,536 unique firms and 189,678 firm-quarter observations, with $\sim$420k workers represented in the sample, which is around half of formal employment in manufacturing. Since firms hire in more than one commuting zone, these correspond to 6,225,319 firm-market-quarter observations, the unit of observation in Tables 1 and 2.

Table 1 shows descriptive statistics. Panel A describes our main estimation sample at the firm-market-quarter level, where we take workers in the UI matched employer-employee dataset at the end of each quarter to obtain total employment for each firm in each market. Unique firm identifiers let us join this dataset with tax records on revenues, capital value, and material expenses, which are used in labor demand estimation. Panel B shows summary statistics for workers in the UI registry at the worker-quarter level.

An important fact is the relationship between within-firm wage dispersion and union status. Note that the distribution of wages is much more compressed in union firms: the within-firm-quarter standard deviation of log wages averages 0.475 at non-union firms versus 0.355 at unionized firms, so dispersion is 33.8% higher at non-union firms. This is consistent with previous evidence on the effect of unions on the within-firm wage distribution. Contracts are such that one base wage (and bonus compensation) is decided for the entire worker body, with small deviations that can be driven by extra hours or specific compensation. This fact is consistent with the assumption in the model that unions and firms bargain over a single (average) wage, which helps the empirical

Table 1. Summary Statistics

	Mean	SD	N
Panel A: Firm-Market-Quarter Level			
Employment (workers)	50.42	105.65	6,225,319
Revenue (million CLP)	19,200	96,800	6,225,319
Unique firms			4,536
Panel B: Worker-Quarter Level			
Wage (CLP)	838,683	708,532	222,610,206
Age (years)	39.26	11.48	222,610,206
Female	0.40	0.49	222,610,206
Formal-sector experience (months)	53.77	43.82	222,610,206

implementation.

We measure employer concentration in each commuting zone with the Herfindahl-Hirschman Index (HHI) of employment shares. Unions are disproportionately located in the most concentrated markets: 56% of the firms that bargain with a union are in the top 20% of commuting zones by HHI. This pattern is consistent with unions potentially acting as a countervailing force to employer market power, but it does not by itself establish that unions form in response to concentration; the pattern is more likely due to firm size and the scope for rent seeking (Azkarate-Askasua and Zerecero, 2025).

In Table 2 we show mean differences between non-union firms and firms that bargain with single- and multi-unit unions, from a simple regression on union-type indicators controlling for time fixed effects and industry-specific linear trends. Unionized firms are considerably larger than non-union firms: employment is about 60% higher at single-unit union firms and 33% higher at multi-unit union firms, and revenue is roughly 8 to 10 times as large for both. Single-unit union firms also hold larger within-industry employment shares, by about 1.2 percentage points, while the difference for multi-unit union firms is small and not statistically significant. Wages are also higher in union firms, as shown previously in other studies: about 21% higher for both union types, with a difference between multi- and single-unit union firms of only half a log point. These wage differences could well be explained by observable and unobservable factors other than collective bargaining through unions. In the model, we explicitly account for both labor supply responses and job characteristics, and for firm characteristics through production function estimation.

Table 2. Mean Differences by Union Status

	(1) Log Employment	(2) Log Revenue	(3) Log Wage	(4) Within-Industry Share
Single-unit union	0.472***	2.260***	0.189***	0.012***
	(0.019)	(0.088)	(0.011)	(0.001)
Multi-unit union	0.286***	2.098***	0.194***	0.004
	(0.029)	(0.093)	(0.014)	(0.003)
Observations	6,225,319	6,225,319	6,225,319	6,225,319
Unique firms	4,536	4,536	4,536	4,536
R-squared	0.054	0.037	0.229	0.003

Notes: Estimates from a firm-market-quarter panel of the estimation sample; observations are firm-market-quarters. Coefficients on indicators for firms bargaining with a single-unit or a multi-unit union; the baseline group is non-union firms. All specifications include quarter fixed effects and industry-specific linear time trends. Standard errors clustered at the market level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

4 Estimation of Labor Demand and Supply

4.1 Empirical Model of Labor Supply

We now describe the empirical implementation of the model in Section 2, starting with labor supply. To derive the labor supply curve $L^S(\cdot)$, we assume a discrete choice model of labor supply with limited availability. Workers indexed by i choose among jobs $j \in \mathcal{J}_m$ in a local labor market m, or an outside option. Since the firm hires randomly out of the supply, a worker who chooses job j is hired with probability ψ_j – the firm's rationing share from Section 4.2 – and receives 0 utility otherwise. The worker therefore chooses between lotteries where one outcome is being hired and the other is getting 0 utility. We assume the expected utility function is $U_{ij} = \psi_j e^{\delta_j + \varepsilon_{ij}}$, with the outside option yielding $e^{\varepsilon_{i0}}$ with probability 1. Indirect utility for worker i from job j is given by:

$$\ln U_{ij} = \delta_j + \ln \psi_j + \varepsilon_{ij} \equiv \underbrace{X_j'\phi + \beta W_j}_{\equiv \delta_j} + \ln(\psi_j) + \varepsilon_{ij}(\lambda) \tag{4}$$

where X_j are observable characteristics of the job over which workers have common preferences, in which we include commuting-zone and year-quarter fixed effects, W_j is the wage offered by job j with sensitivity β, and $\varepsilon_{ij}(\lambda)$ is an idiosyncratic taste shock following a nested logit (GEV) distribution, with nests given by the industry of job j within the market: this makes jobs in the same industry closer substitutes for one another than jobs in different industries.

A modeling choice worth flagging is how $\ln(\psi_j)$ enters utility: it is deterministic from the point of view of the worker. Because of this, and because it depends on demand characteristics, it acts as a shifter of preferences for jobs: if a job is less available because ψ_j is small, its utility is penalized. Because ψ_j depends on demand factors (as explained in more detail below when we solve for ψ_j in Section 4.2.1) we are able to separately identify it from δ_j.

$$s_j = \underbrace{\frac{\psi_j^{1/\lambda}\exp\left(\frac{\delta_j}{\lambda}\right)}{\sum_{k\in\mathcal{J}_{g(j),m}}\psi_k^{1/\lambda}\exp\left(\frac{\delta_k}{\lambda}\right)}}_{s_{j|g(j)},\text{ probability of picking } j \text{ within its industry}} \cdot \underbrace{\frac{\left(\sum_{k\in\mathcal{J}_{g(j),m}}\psi_k^{1/\lambda}\exp\left(\frac{\delta_k}{\lambda}\right)\right)^{\lambda}}{1+\sum_{g\in\mathcal{G}_m}\left(\sum_{k\in\mathcal{J}_{g,m}}\psi_k^{1/\lambda}\exp\left(\frac{\delta_k}{\lambda}\right)\right)^{\lambda}}}_{s_{g(j)},\text{ probability of picking industry } g} \tag{5}$$

where $g(j)$ is the industry of job j and $\mathcal{G}_m$ is the set of industries in market m. This corresponds to the worker deciding which industry g to work in, and then which firm in that industry. The first term is the probability of picking job j conditional on its industry; the second is the probability of picking that industry in the first place. Given the distributional assumptions, the nesting parameter λ governs within-industry substitution: as $\lambda \to 0$, jobs in the same industry become near-perfect substitutes, while as $\lambda \to 1$, the nested model collapses to a standard logit. Finally, note how the rationing rates ψ_j enter here: when a job becomes more (less) available, the probability that workers apply to it is higher (lower). Likewise, as other jobs become more available, they become relatively more attractive, since workers are more likely to be hired there, and vice versa.

The 1 in the denominator is the outside option, whose mean utility is normalized to zero. Labor supply to job j, the number of workers who apply, and the resulting employment are

$$L_j^S = \bar{L}\, s_j, \qquad L_j = \psi_j L_j^S = \psi_j\, \bar{L}\, s_j, \tag{6}$$

where $\bar{L}$ is the size of the local labor market (Section 3). Employment is consistent with the firm's problem in Section 2.

Labor supply elasticities. The rationing mechanism gives this model a qualitatively sharper prediction than a standard logit: a firm's labor supply elasticity depends on which side of the

market it is on. A demand-constrained firm ($\psi_j < 1$) has employment pinned exactly at its labor demand $L_j^D(W_j)$ by construction, so its employment does not respond to its own wage through the supply side at all – it has no monopsony power in this regime, only the pass-through of a standard labor demand elasticity. A supply-constrained firm ($\psi_j = 1$) instead faces the usual nested-logit own- and cross-wage elasticities – positive own-wage, negative cross-wage, governed by λ and market concentration as before – but now transmitted through the shared within-industry and across-industry sums in equation (5) rather than directly. Multi-unit unions internalize the cross-firm term in the supply-constrained regime when bargaining a wage increase at one of their firms, as captured by the diversion term in equation (2).

4.1.1 Estimation and Identification

Because ψ_j enters the shares that the estimating equation below is built on, it has to be solved for in equilibrium using the hiring game of Section 4.2.1 before the equation can be formed.

The parameters to be estimated are the coefficients on the observable job characteristics and fixed effects, ϕ, the wage sensitivity β, and the nesting parameter λ; call this vector θ. Following Berry (1994), the nested logit shares in equation (5) can be inverted in closed form. Let s_{0m} be the share of the outside option in market m and $s_{j|g(j)}$ the within-industry share of job j. Taking logs of the ratio s_j/s_{0m} and substituting out the industry sum gives the linear estimating equation

$$\ln s_j - \ln s_{0m} - \ln \psi_j = X_j'\phi + \beta W_j + (1-\lambda)\ln s_{j|g(j)} + \eta_j, \tag{7}$$

where η_j is an error term that captures determinants of the observed shares not included in the model. The shares in equation (7) are choice probabilities. Because only a fraction ψ_j of the workers who choose job j are hired, we recover them from observed employment as $s_j = L_j/(\psi_j \bar{L})$. The rationing share therefore enters the estimating equation twice: it corrects the observed share for applicants who were not hired, and it appears as the availability term $\ln \psi_j$, which moves to the left-hand side because its coefficient is fixed at one by equation (4). For supply-constrained firms ($\psi_j = 1$), which are the large majority of firms and include every non-union firm by construction, the equation reduces to the standard nested logit inversion. Since ψ_j depends in turn on (δ_j, λ) through the hiring game, we estimate θ and the rationing shares jointly: given a candidate θ we

solve the hiring game for ψ in every market, form the left-hand side of equation (7), estimate θ by two-stage least squares, and iterate until θ and ψ are mutually consistent. We refer the reader to Appendix B for details.

Equation (7) has two endogenous regressors: the wage W_j and the within-industry share $\ln s_{j|g(j)}$, both of which are correlated with the error term η_j.

Wages are set in equilibrium, so they are correlated with unobserved determinants of labor supply in η_j. To address this simultaneity problem, we need instruments Z_j that shift demand without shifting supply. One way of achieving this is to use instruments that affect the wage policies of the firm but are uncorrelated with worker characteristics in a particular market. We use the demographics of workers in the *other markets* where the firm is posting jobs. The idea is that firms usually make an effort to reduce cross-market wage dispersion (Hjort et al., 2026; Hazell et al., 2022), which leads to determinants of wages in other markets being correlated with wages in the market of interest, but excluded from said market's unobservable characteristics. Another way to explain the intuition is that competitive conditions in the other labor markets where the firm is posting affect the wage policy of the firm. This will be reflected in the wages in the market of interest, but is not correlated with unobserved determinants of labor supply in that market.

To identify the nesting parameter λ, we instrument the within-industry share with the number of other firms in the same industry and market, excluding firm j itself. Conditional on the market fixed effects, a job in a more crowded industry has a mechanically smaller within-industry share, and the number of competitors is determined by entry decisions made before the current period's shocks to labor supply are realized. How much the within-industry share falls with the number of competitors identifies how closely workers substitute across jobs in the same industry, which is the role of λ. For this instrument to have variation, X_j includes market fixed effects but not industry fixed effects, since industry fixed effects within a market would absorb the firm count entirely.

Table 3 shows the parameter estimates of the labor supply model. We find that our results are consistent with the labor literature. The OLS wage coefficient is small and negative, which would imply that workers prefer lower-paying jobs. This is what simultaneity between wages and unobserved determinants of labor supply predicts: jobs that attract workers for reasons we do not observe can pay lower wages, so the wage is negatively correlated with the error term and OLS is biased downward. Instrumenting the wage reverses the sign and raises the coefficient by an

Table 3. Labor Supply Parameter Estimates

	OLS	IV, Nested Logit
$\hat{\beta}$ (wage sensitivity)	-0.0574^{*}	0.4025^{***}
	(0.0312)	(0.0162)
$\hat{\lambda}$ (nesting parameter)	—	0.3459^{***}
		(0.0001)
Markets (commuting zone $\times$ quarter)	256	256
Firms	4,536	4,536
Observations	189,678	189,678
First-stage F-statistic (Kleibergen-Paap)	—	23.4

Notes: Standard errors clustered by market in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Two-stage least squares estimates of equation (7), iterated with the hiring-game equilibrium. The wage is instrumented with the demographics of workers in other markets where the firm is posting jobs, and the within-industry share with the number of other firms in the same industry and market; see Section 4.1.1. All specifications include market fixed effects. The OLS column does not instrument the wage.

order of magnitude. The IV wage coefficient is positive and precisely estimated, and the nesting parameter λ is statistically significant and closer to zero than to one, indicating that within-industry substitution across jobs is important, again consistent with recent studies of workers' response to labor market conditions (Arnold, 2019; Schubert et al., 2026).

Figure 3 shows the distribution of implied firm-level own-wage labor supply elasticities. We find that the median firm-level own-wage elasticity is 3.16. These results are in the range found in other studies in similar settings (Berger et al., 2022; Sharma, 2025; Felix, 2026).

4.2 Labor Demand

To describe labor demand behavior, we start from the firm's profit maximization problem. We treat the bargaining problem as short-run: capital K_j is chosen before wages are bargained and its cost is sunk, while labor and materials are flexible. As in Section 2 and following Azkarate-Askasua and Zerecero (2025), the firm does not choose employment directly. Given a wage W_j, it chooses the rationing rate ψ_j, the share of the workers who apply at W_j that it hires:

$$\max_{\psi_j \in [0,1]} \pi_j(\psi_j, W_j) = R_j\big(\psi_j L_j^S(W_j)\big) - W_j\, \psi_j L_j^S(W_j), \tag{8}$$

where $R_j(L_j) = F(K_j, L_j)\exp(\omega_j)$ is the firm's (value-added) revenue, $F(\cdot)$ is the production technology, and ω_j is a Hicks-neutral productivity shock observed by the firm but not by the

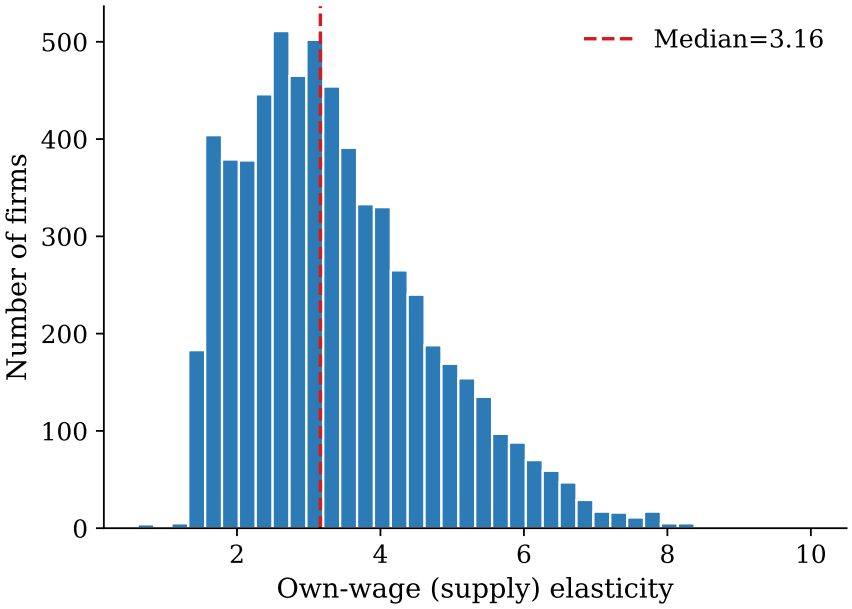


Figure 3. Implied Firm-Level Labor Supply Elasticities

Notes: Distribution of the firm-level own-wage labor supply elasticity across supply-constrained firms ($\psi_j = 1$), evaluated at equilibrium employment.

econometrician. The solution is the rationing condition in equation (1). If the firm's unconstrained labor demand $L_j^D(W_j)$, defined by $R_j'(L_j^D(W_j)) = W_j$, falls short of the number of workers who apply, it rations to $\psi_j = L_j^D(W_j)/L_j^S(W_j) < 1$ and employment lies on the labor demand curve. Otherwise it hires everyone who applies, $\psi_j = 1$, and employment equals labor supply. Two objects from this problem enter the bargaining model. The first is the output elasticity of labor $\beta_j^L \equiv \partial \ln R_j / \partial \ln L_j$, which gives the marginal revenue product $R_j'(L_j) = \beta_j^L R_j / L_j$ and bounds the union markup at $1/\beta_j^L$ in equation (3). The second is the labor demand elasticity,

$$\varepsilon_j^D \equiv \frac{d \ln L_j^D}{d \ln W_j} = \left(\beta_j^L - 1 + \frac{\partial \ln \beta_j^L}{\partial \ln L_j} \right)^{-1}, \tag{9}$$

which determines how much employment a firm gives up once a bargained wage pushes it into rationing.

4.2.1 The Hiring Process

Our approach is related to the rationing mechanism in Angerhofer et al. (2026), who model offers at the individual level, made sequentially to the nearest applicant. We instead take an aggregate

approach suited to our discrete-choice labor supply system, in which workers are hired randomly, making the probability of getting hired equal to the firm's rationing rate ψ_j.

Firm j receives $\bar{L}s_j$ applicants and hires $\psi_j \bar{L}s_j$ of them. Given the other firms' rationing rates ψ_{-j}, its employment is strictly increasing in its own ψ_j, since hiring a larger share of applicants also attracts more of them. The firm's problem in equation (8) then has a simple solution. If it can, the firm sets ψ_j so that it hires exactly its labor demand $L_j^D(W_j)$; otherwise it hires every applicant, setting $\psi_j = 1$. The best response function, which summarizes the solution to the profit maximization problem, is therefore

$$\psi_j(\psi_{-j}) = \min\left\{1, \left[\frac{L_j^D(W_j)}{L_j^S(\cdot)}\right]\right\}. \tag{10}$$

Proposition 1. *For given wages $\{W_j\}_{j\in\mathcal{J}_m}$, the hiring game in market m has a unique Nash equilibrium $\psi^* = (\psi_1^*, \ldots, \psi_J^*)$.*

The formal proof is in Appendix A; here we give the intuition. First, rationing rates are strategic complements: if other firms increase the share of applicants they hire, their jobs become more attractive and firm j's labor supply falls, so firm j has to hire a larger share of its own applicants to keep employment at its labor demand. Since each firm's strategy set $[0, 1]$ is a lattice and its best response is increasing in the other firms' strategies, existence follows from the monotone-games result of Milgrom and Shannon (1994). Uniqueness follows from reducing the equilibrium conditions to a single scalar fixed point in the market-level aggregate of the labor supply system, which has exactly one solution. The same reduction delivers a fast solution algorithm, which helps with the computation of estimates and counterfactuals.

4.2.2 Solution Algorithm

To summarize the estimation and solution procedure: the production function is estimated first, from the tax records alone (Section 4.2.3), so each firm's labor demand $L_j^D(W_j)$ at its observed wage is known before we turn to labor supply. Given labor demand, we recover the rationing rates and the supply parameters θ jointly, in four steps. First, we set $\psi_j = 1$ for every firm without a union, since these firms set the oligopsony wage, which lies on the supply-constrained segment of the bargaining locus, and we start from $\psi_j = 1$ for unionized firms as well. Second, given the rationing

rates, we recover choice probabilities from observed employment as $s_j = L_j/(\psi_j \bar{L})$ and estimate θ by two-stage least squares on the inverted shares (equation (7)), which also delivers the mean utilities δ_j. Third, given δ_j and λ, we solve the hiring game market by market for the equilibrium rationing rates of unionized firms, which is unique by Proposition 1. Fourth, we return to the second step with the new rationing rates and iterate until θ and ψ stop changing. Labor demand is what separates rationing from preferences. A unionized job with few workers relative to the applicants its mean utility would attract is a job that rations, because the firm's labor demand, which comes from the production side and not from workers' preferences, falls short of its applicant pool. Appendix B gives the algorithm in full.

4.2.3 Production Function Estimation

We specify $F(\cdot)$ as a translog in capital and labor. We work with value added under the assumption that gross output is Leontief in materials, which is consistent with the manufacturing industries in our sample. The translog lets both elasticities above vary across firms with their input mix. This flexibility matters for the bargaining model. Under Cobb-Douglas, β_j^L is constant along the demand curve, so equation (9) reduces to $\varepsilon_j^D = 1/(\beta_j^L - 1) < -1$, and labor demand is elastic at every firm by construction. A union that maximizes the wage bill would then never bargain a firm into rationing, since a higher wage always lowers $W_j L_j$ on the demand curve. That would rule out the demand-constrained segment of the bargaining locus in Figure 1 by functional form instead of by the data.

We estimate the production function following Ackerberg et al. (2015) (ACF), using materials as a proxy for the unobserved productivity shock ω_j. The timing assumptions underlying the ACF approach are consistent with the bargaining model: capital is fixed one period before ω_j is realized, while labor and materials are chosen after, so the labor input this procedure recovers is the one relevant for the marginal-revenue-product objects entering the firm's bargaining problem. Identification uses lagged labor as an instrument, which is valid as long as innovations to productivity are realized after the previous period's labor and capital decisions and are therefore uncorrelated with these lagged choices. This method has been applied in closely related labor settings to recover marginal revenue products under imperfect competition (Berger et al., 2022; Kroft et al., 2025); we

refer the reader to Ackerberg et al. (2015) and to Appendix C for the full estimation details and the estimated coefficients (Table A1).

4.3 Demand-Side Elasticities, Rationing, and Unionization

Table 4. Labor Output and Labor Demand Elasticities

	Output elasticity ($\hat{\beta}_j^L$)			Demand elasticity ($\hat{\varepsilon}_j^D$)		
	Mean	Median	Range	Mean	Median	Range
All firms	0.642	0.647	[0.341, 0.916]	−2.399	−1.988	[−5.275, −0.823]
Single-unit unions	0.565	0.560	[0.315, 0.854]	−1.917	−1.631	[−3.857, −1.007]
Multi-unit unions	0.349	0.316	[0.272, 0.472]	−1.087	−1.010	[−1.364, −0.823]

Notes: Firm-level output elasticity of value added with respect to labor, and labor demand elasticity (equation (9)), evaluated at each firm's equilibrium employment. The production function is a translog in capital and labor estimated via Ackerberg et al. (2015) with materials as the productivity proxy; see Appendix C. Range: 5th to 95th percentile.

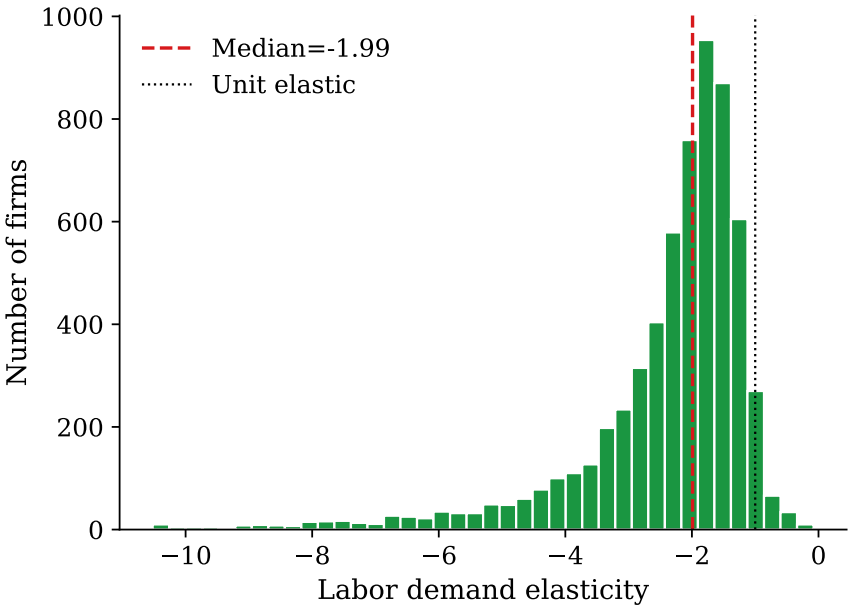


Figure 4. Distribution of Labor Demand Elasticities

Notes: Distribution of the firm-level labor demand elasticity $\hat{\varepsilon}_j^D$ (equation (9)) across supply-constrained firms ($\psi_j = 1$), evaluated at equilibrium employment.

Table 4 and Figure 4 report the results. Output elasticities of labor are well below one, with a median of 0.647, and are lower at union firms. This sets the size of the union's potential markup over $R_j'(L_j)$: in the limit where the union captures the full bargaining weight ($\omega_j^u \to 1$), the wedge in equation (3) becomes $\lambda_j \to 1/\beta_j^L$, so a lower β_j^L leaves more rents for the union to extract. Labor

demand is elastic at the median firm (-1.99) and inelastic at only a small share of firms. Consistent with this, few unionized firms are bargained into rationing, and those that are have demand close to unit-elastic (-1.01 at the median). Interestingly, the rationed firms are almost exclusively those bargaining with multi-unit unions.

5 Effect of Multi-Unit Bargaining on Labor Outcomes

We now proceed to quantify the effect of multi-unit bargaining on wages, employment, and worker welfare. In order to do so, we use our estimates together with the bargaining first-order condition (3) to recover the bargaining power parameters γ_j^u for each bargaining pair. We then use these parameters to compute counterfactual wages and employment under the assumption that all bargaining units are single-unit unions, and compare these counterfactuals with the observed outcomes. We proceed in four steps. First, we describe the union objective and how we recover the bargaining power parameters. Second, we compute counterfactual wages and employment under the assumption that all bargaining units are single-unit unions, and compare these counterfactuals with the observed outcomes. Third, we explore heterogeneity in bargaining power and the diversion externality across firms. Fourth, we check how well the underlying labor supply and demand system fits observed wages for both supply- and demand-constrained union firms, independent of the bargaining layer built on top of it.

5.1 Union Objective and Implied Bargaining Power

The union objective. We assume that the union's objective is the total wage bill of the jobs it represents,

$$F_u(\mathbf{W}^u, \mathbf{L}^u) = \sum_{k \in \mathcal{B}_u} W_k L_k. \tag{11}$$

This objective can be justified in three ways. First, it describes a "bureaucratic" union that cares about its own survival and resources. Union dues are deducted from covered workers' wages, and in Chile workers who receive the benefits of an agreement through benefit extension also pay dues (Section 3). Total dues are therefore proportional to the wage bill of the bargaining unit, and they are what allow the union to survive and to fund its activities. Second, the wage bill is the

classic objective of a rent-maximizing union (Dunlop, 1944), and it is what the utilitarian union of McDonald and Solow (1981) maximizes when workers are risk neutral and their outside option is normalized to zero. In logs, it gives equal weight to the wage level and to employment at each firm, like an equal-weighted Cobb-Douglas utility over the two. Third, the employment term is what makes the diversion mechanism of Section 2 possible. A union that cared only about wages would not internalize that a higher wage at one of its firms moves workers away from its other firms. Under the wage bill, that diversion lowers the union's objective one-for-one with the wage bill lost at the other firms.

Each union u then bargains with every firm $j \in \mathcal{B}_u$ separately and simultaneously, taking the wages agreed at its other firms, W_k^* for $k \neq j$, as given. The bargained wage solves

$$W_j^* = \arg\max_{W_j} \left[\pi_j(\psi_j, W_j) - \bar{\pi}_j\right]^{1-\gamma_j^u} \times \Big[\underbrace{W_j L_j(W_j, \mathbf{W}_{-j}^*) + \sum_{k \in \mathcal{B}_u, k \neq j} W_k^* L_k(W_j, \mathbf{W}_{-j}^*)}_{F_u \text{ given the other agreements}} - \underbrace{\sum_{k \in \mathcal{B}_u, k \neq j} W_k^* L_k(W_j^*, \mathbf{W}_{-j}^*)}_{d_{uj}}\Big]^{\gamma_j^u}, \tag{12}$$

subject to the firm's right to manage: after the wage is agreed, the firm sets its rationing rate ψ_j unilaterally according to equation (1), so that employment is $L_j = \psi_j L_j^S(W_j)$ and W_j never exceeds the marginal revenue product of labor $R_j'(L_j)$. The firm's gain from trade is its profit $\pi_j(\psi_j, W_j) = R_j(\psi_j L_j^S(W_j)) - W_j \psi_j L_j^S(W_j)$, evaluated at this rationing rate, relative to its disagreement profit $\bar{\pi}_j$. Since capital is sunk in the short run and firms cannot replace workers during a bargaining dispute in Chile (Section 3), a firm that fails to reach an agreement produces nothing, and we set $\bar{\pi}_j = 0$. The union's disagreement point d_{uj} is the wage bill it earns at its other firms in the unit. Its gain from trade at the agreement is therefore the wage bill at firm j, $W_j^* L_j$, plus the change in the wage bill at its other firms caused by moving W_j. For a single-unit union ($\mathcal{B}_u = \{j\}$) the sums are empty and the union's surplus is just $W_j L_j$.

Taking the first-order condition of equation (12) delivers the wage equation (3) at a supply-constrained firm, and a condition that pins down the wage along the labor demand curve at a demand-constrained one (both are derived in Appendix D). Under the wage bill objective, the elasticity of the union's gain from trade with respect to the wage, ϕ^u in equation (2), takes the

simple form

$$\phi_j^u = 1 + \epsilon^{L_j,W_j} + \sum_{k \in \mathcal{B}_u, k \neq j} \frac{W_k L_k}{W_j L_j} \epsilon^{k,j}, \qquad \epsilon^{L_j,W_j} = \begin{cases} e_j & \text{if } \psi_j = 1, \\ \varepsilon_j^D & \text{if } \psi_j < 1, \end{cases} \tag{13}$$

where $\epsilon^{k,j}$ is the elasticity of employment at firm k with respect to the wage at j, evaluated in firm j's regime. The 1 is the own-wage effect: at fixed employment, a 1% higher wage raises the wage bill at j by 1%. The second term is the own-employment effect, which raises the wage bill along the labor supply curve and lowers it along the labor demand curve. The sum is the diversion effect, weighted by the size of each other firm's wage bill relative to j's. For a single-unit union the sum is empty, so $\phi_j^u = 1 + e_j > 0$ when the firm is supply-constrained and $\phi_j^u = 1 + \varepsilon_j^D$ when it rations, which is positive only if labor demand is inelastic.

In the supply-constrained regime the cross-elasticities are the labor supply cross-elasticities $e_{kj} < 0$ (Appendix B), and we can write $\phi_j^u = (1 + e_j)(1 + \kappa_j)$, where $\kappa_j \leq 0$ is the diversion term

$$\kappa_j \equiv \frac{1}{1 + e_j} \sum_{k \in \mathcal{B}_u, k \neq j} \frac{W_k L_k}{W_j L_j} e_{kj}, \tag{14}$$

obtained by factoring $1 + e_j$ out of the diversion sum in equation (13); since $e_{kj} < 0$, $\kappa_j \leq 0$, with equality only for a single-unit union. The bargaining weight is then

$$\omega_j^u = \frac{\gamma_j^u (1 + \kappa_j)}{\gamma_j^u (1 + \kappa_j) + 1 - \gamma_j^u}. \tag{15}$$

For a single-unit union, $\kappa_j = 0$ and the bargaining weight equals the bargaining power, $\omega_j^u = \gamma_j^u$. For a multi-unit union, diversion ($\kappa_j < 0$) pushes ω_j^u below γ_j^u, so the same bargaining power buys a smaller weight on the union markup.

With labor supply and demand estimated, every object in the first-order conditions except γ_j^u is either observed or estimated, so bargaining power can be recovered pair by pair in closed form. The formula depends on the regime of the firm. Note that we are explicitly controlling for equilibrium objects in the equation, with bargaining power identified from the residual wages conditional on supply and demand decisions and heterogeneity in firm and market characteristics. For a supply-constrained firm, given the observed wage W_j, the marginal revenue product $R_j'(L_j)$

and output elasticity β_j^L (Section 4.2), and the markdown $\mu_j = e_j/(1+e_j)$ (Section 4.1), the wage equation identifies the bargaining weight in closed form,

$$\hat{\omega}_j^u = \frac{W_j/R_j'(L_j) - \mu_j}{1/\beta_j^L - \mu_j}, \tag{16}$$

which is the position of the observed wage between the oligopsony wage ($\omega_j^u = 0$) and the wage at the union's maximal markup ($\omega_j^u = 1$). Inverting equation (15) then gives bargaining power,

$$\hat{\gamma}_j^u = \frac{\hat{\omega}_j^u}{\hat{\omega}_j^u + (1-\hat{\omega}_j^u)(1+\kappa_j)}, \tag{17}$$

which equals $\hat{\omega}_j^u$ for single-unit pairs. For multi-unit pairs it exceeds $\hat{\omega}_j^u$ by an amount governed by the estimated cross-elasticities in κ_j. For a rationed firm ($\psi_j < 1$), the observed wage equals the marginal revenue product, and the first-order condition in this regime (equation (D.2)) gives bargaining power directly,

$$\hat{\gamma}_j^u = \frac{b_j}{\phi_j^u + b_j}, \qquad b_j \equiv \frac{\beta_j^L}{1-\beta_j^L}, \tag{18}$$

with ϕ_j^u from equation (13) evaluated at the labor demand elasticity $\hat{\varepsilon}_j^D$ (Section 4.2). Rationing is consistent with the bargaining model only if $\phi_j^u > 0$; for single-unit pairs this is the requirement that $\hat{\varepsilon}_j^D > -1$. The only pairs for which bargaining power is not point-identified are those whose wage is exactly the competitive wage: any γ_j^u in an interval (equation (D.3)) rationalizes that wage, so for these pairs we recover that interval. No numerical root-finding is needed in any of the three cases. Details are in Appendix D.

In Figure 5 we show the distribution of implied bargaining power $\hat{\gamma}_j^u$ across bargaining pairs, separately for single- and multi-employer bargaining units. Multi-unit unions have higher implied bargaining power than single-unit unions: mean $\hat{\gamma}^{\text{single}} = 0.23$ versus mean $\hat{\gamma}^{\text{multi}} = 0.31$. This is exactly what the mechanism in Section 2 predicts: because a multi-unit union internalizes the employment-diversion externality in equation (2), its objective function is less elastic to the wage at any single firm than a single-unit union's, all else equal, which on its own would generate a *smaller* bargaining weight ω_j^u (equation (3)) and thus a lower bargained wage at that firm. Since bargained wages for multi-unit pairs in the data are not correspondingly lower, our procedure rationalizes part of the observed outcomes by attributing higher underlying bargaining power γ_j^u to multi-unit

unions: multi-unit unions appear to need, and to have, more raw bargaining strength to offset the diversion externality they internalize. There is also substantial heterogeneity within each group, which is important for understanding the scope for market power on the union side and, in turn, the potential for unions to extract rents from firms.

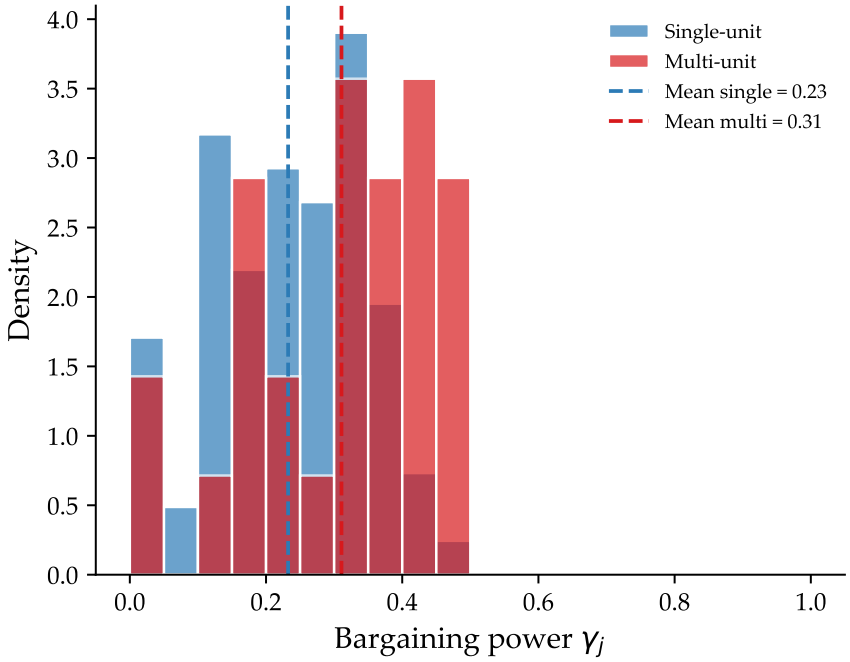


Figure 5. Implied Bargaining Power, Single- vs. Multi-Unit Unions

Notes: Distribution of the implied bargaining power parameter $\hat{\gamma}_j^u$ recovered from the bargaining first-order condition (equation (3)), separately for single-employer and multi-employer bargaining units. See Appendix D for the recovery procedure.

5.2 Model Fit: Supply- and Demand-Constrained Union Firms

We test the fit of the underlying labor supply and demand system to observed wages for both supply- and demand-constrained union firms, independent of the bargaining layer built on top of it. We compare observed and model-predicted wages. The results are shown in Figure 6. The dashed line is the 45° line in both panels. For the 60 supply-constrained firms (Panel a), predicted and observed wages are highly correlated ($r = 0.93$, $R^2 = 0.87$), with most points lying close to the 45° line. For the 78 demand-constrained firms (Panel b), the fit is weak ($r = 0.20$, $R^2 = 0.04$): predicted wages are only loosely related to observed wages and sit visibly below the 45° line for most firms.

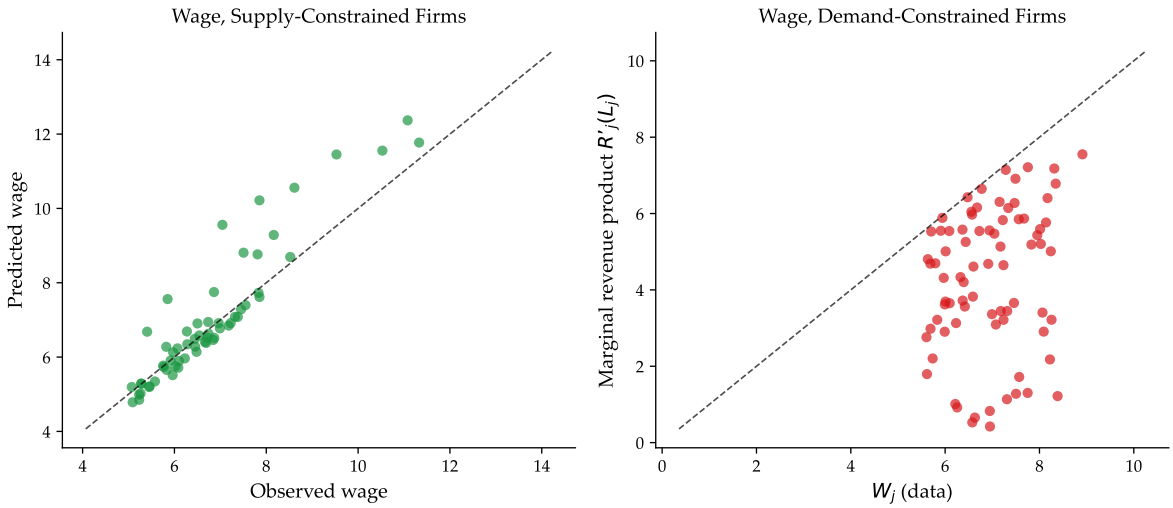


Figure 6. Model Fit, Supply- and Demand-Constrained Union Firms

Notes: Panel (a): supply-constrained union firms ($\psi_j = 1$) – the oligopsony wage W_j^O, recovered from the labor supply and demand system alone, against the observed wage W_j. Panel (b): demand-constrained union firms ($\psi_j < 1$) – the marginal revenue product $R'_j(L_j)$, evaluated at observed employment but never fit to the observed wage, against W_j. The dashed line is the 45° line in both panels.

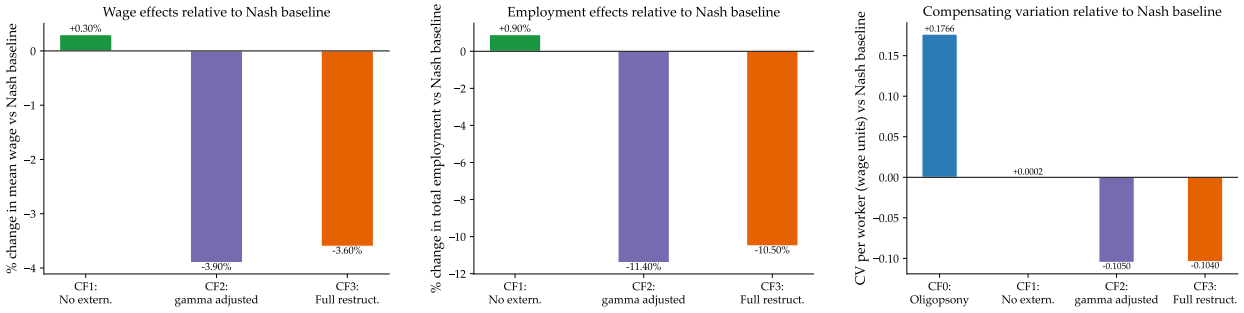
5.3 Effects of Bargaining Power Gains and Employment Diversion in the Multi-Unit Case

To address the labor outcome effects of multi-unit union bargaining, we compute counterfactual wages and employment for the multi-unit unions in three different ways. First, we relocate their distribution of bargaining power parameters γ_j^u such that they have the same mean bargaining power as single-unit unions. Then we solve the new Nash-in-Nash equilibrium with the new bargaining power parameters and compute the new wages and employment levels. We then compare these counterfactual outcomes with the observed outcomes to quantify the effect of multi-unit bargaining on wages and employment.

We decompose this effect into two channels: the employment-diversion externality that a multi-unit union internalizes (CF1: multi-unit pairs converted to single-unit, γ_j^u held fixed), and the bargaining power gap documented in Section 5.1 (CF2: γ_j^u for multi-unit pairs shifted toward the single-unit mean by $\Delta = (\bar{\gamma}^{\text{single}} - \bar{\gamma}^{\text{multi}})/\bar{\gamma}^{\text{multi}}$, cross-firm structure held fixed). CF3 combines both channels. Figure 7 reports the resulting change in the mean wage (Panel a) and total employment (Panel b) of union firms, and worker welfare (Panel c), all relative to the observed (Nash) baseline and restricted to the same three scenarios throughout. Isolated from the bargaining-power

gap, the externality channel alone (CF1: +0.3% wage) is mildly wage-*increasing*: once a multi-unit union no longer internalizes the diversion of employment to its other firms, its objective is more elastic to the wage at any one firm, so it settles for a slightly higher wage there. Equalizing bargaining power alone (CF2: −3.9%) is the dominant force, lowering wages substantially more than the combined effect on its own. Combining both channels (CF3: −3.6%) sits between CF1 and CF2 – the externality channel partially offsets the bargaining-power channel, pulling the combined effect back part of the way toward zero.

Employment moves in the same direction as wages (CF1: +0.9%; CF2: −11.4%; CF3: −10.5% – total employment of union firms relative to the Nash baseline), consistent with the labor supply elasticities each firm faces. Worker welfare (Panel c) tracks the same ordering: compensating variation (CV) per worker is small and positive (+0.0002) under the externality-only scenario, and substantially negative under both the bargaining-power-only scenario (−105,000 CLP) and the combined scenario (−104,000 CLP).



(a) Wages **(b)** Employment **(c)** Worker welfare (CV)

Figure 7. Wage, Employment, and Welfare Effects of Multi-Unit Bargaining Power, by Channel

Notes: Percent change in the mean wage (a) and mean employment (b) of union firms, and compensating variation per worker (c), relative to the observed Nash-in-Nash baseline, under the same three counterfactual scenarios in every panel: CF1 removes the multi-unit externality only (firms reassigned to single-unit bargaining, $\hat{\gamma}_j^u$ unchanged); CF2 equalizes bargaining power only (multi-unit $\hat{\gamma}_j^u$ shifted to the single-unit mean, cross-firm structure unchanged); CF3 combines both. See Appendix D for the recovery of $\hat{\gamma}_j^u$.

5.4 Heterogeneity in Bargaining Power and the Diversion Externality

The decomposition in Section 5.3 established that the bargaining-power gap $\hat{\gamma}_j^u$ documented in Section 5.1, not the diversion externality on its own, is driving a multi-unit wage premium. This leaves open *why* multi-unit unions appear to hold more raw bargaining power, and what the

diversion externality looks like once we stop averaging over bargaining pairs. We examine three things: how $\hat{\gamma}_j^u$ correlates with firm productivity and size differently for single- and multi-unit unions; how the strength of the diversion externality varies with the characteristics of the two firms in a pair; and how the resulting wage and employment changes are distributed across firms, placing every counterfactual on the same oligopsony-to-competitive spectrum as Figure 1.

Table 5 reports correlations between recovered bargaining power $\hat{\gamma}_j^u$ and employer size and productivity, both unweighted and employment-weighted. We report both because a handful of large employers can dominate a weighted correlation, and we additionally drop the five largest firms by employment as an outlier-sensitivity check. The table shows weak, sign-inconsistent correlations for single-unit pairs throughout (essentially zero against both size and productivity, in every column). Multi-unit pairs show a larger but fragile pattern instead: the correlation with size is negative in every specification but grows sharply once employment-weighted (-0.094 unweighted versus -0.671 employment-weighted), and the correlation with productivity is sizable unweighted ($+0.436$) but nearly disappears once employment-weighted (-0.116). Both point the same way – whatever relationship exists for multi-unit pairs is driven by a handful of large employers rather than a stable feature of the data – and the pooled (all-union) correlation with size even flips sign between the unweighted ($+0.185$) and employment-weighted (-0.149) columns. This suggests that our results are not driven by a systematic sorting of multi-unit unions into larger or more productive firms, and that the observed heterogeneity in bargaining power is not mechanically related to firm size or productivity.

Table 5. Correlation Between Recovered Bargaining Power $\hat{\gamma}_j^u$ and Firm Characteristics

	N	Unweighted	Employment-weighted	Outlier-trimmed
Panel A: vs. log *employment*				
Single-unit	190	+0.026	+0.116	+0.007
Multi-unit	63	-0.094	-0.671	-0.112
All union	253	+0.185	-0.149	+0.158
Panel B: vs. productivity ω_j				
Single-unit	190	-0.008	-0.061	+0.013
Multi-unit	63	+0.436	-0.116	+0.464
All union	253	+0.088	+0.017	+0.088

Notes: Correlation between the recovered bargaining-power parameter $\hat{\gamma}_j^u$ and (Panel A) log employment, (Panel B) productivity ω_j, separately for single- and multi-unit bargaining pairs and pooled. "Outlier-trimmed" excludes the five largest firms by employment. See Appendix D.

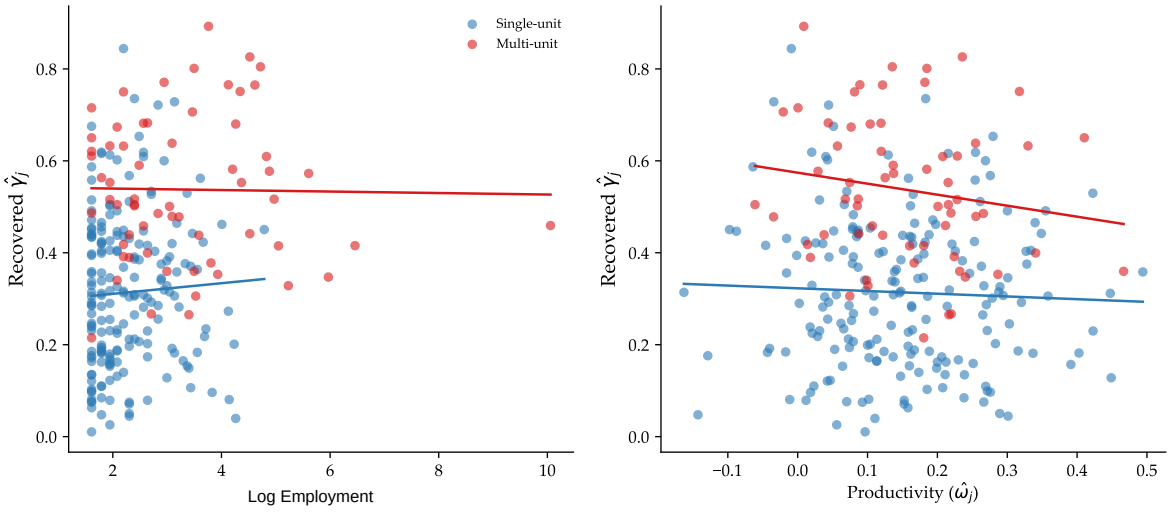


Figure 8. Recovered Bargaining Power vs. Firm Size and Productivity

Notes: Scatter of $\hat{\gamma}_j^u$ against $\log L_j$ (left) and ω_j (right), with a fitted line for single-unit and multi-unit pairs separately.

We also study the correlation between the diversion effect and the characteristics of the bargaining pair. We find no significant correlation (almost zero in every case) between the diversion effect and the size of the two firms in a pair, nor with their productivity or the distance between them. This is consistent with the fact that the diversion effect is a function of the cross-elasticities of labor supply, which depend on substitution patterns across jobs rather than on the size or productivity of the firms in a pair. We also find no correlation between the diversion effect and the bargaining power of the union at either firm in a pair. This suggests that the diversion effect is not systematically related to firm size, productivity, or union bargaining power in a mechanical way, consistent with our model.

In Figure 9, we show the comparison of wages and employment in equilibrium versus perfect competition, and versus the oligopsonistic solution. This exercise lets us see *which* firms move furthest from each benchmark depending on union type. Relative to perfect competition, almost every firm sits below the 45° line in employment. In wages, we observe a subset of demand-constrained multi-unit pairs (outlined markers, bottom-left panel), which sit *above* the 45° line – these firms are already paid at their own labor-demand curve rather than a monopsony-discounted wage, so the market-wide competitive wage can fall below their firm-specific demand-determined wage even as competitive employment, which is governed by the economy-wide labor-supply response rather than any one firm's demand curve, remains higher for them too (bottom-right

panel). Some firms have an observed wage lower than their oligopsony wage. This happens because in the oligopsony regime no firm rations, which might increase the availability of jobs everywhere else in the market, making it more competitive. Angerhofer et al. (2026) find a very similar phenomenon for public schools in the US, where the oligopsonist is the school district.

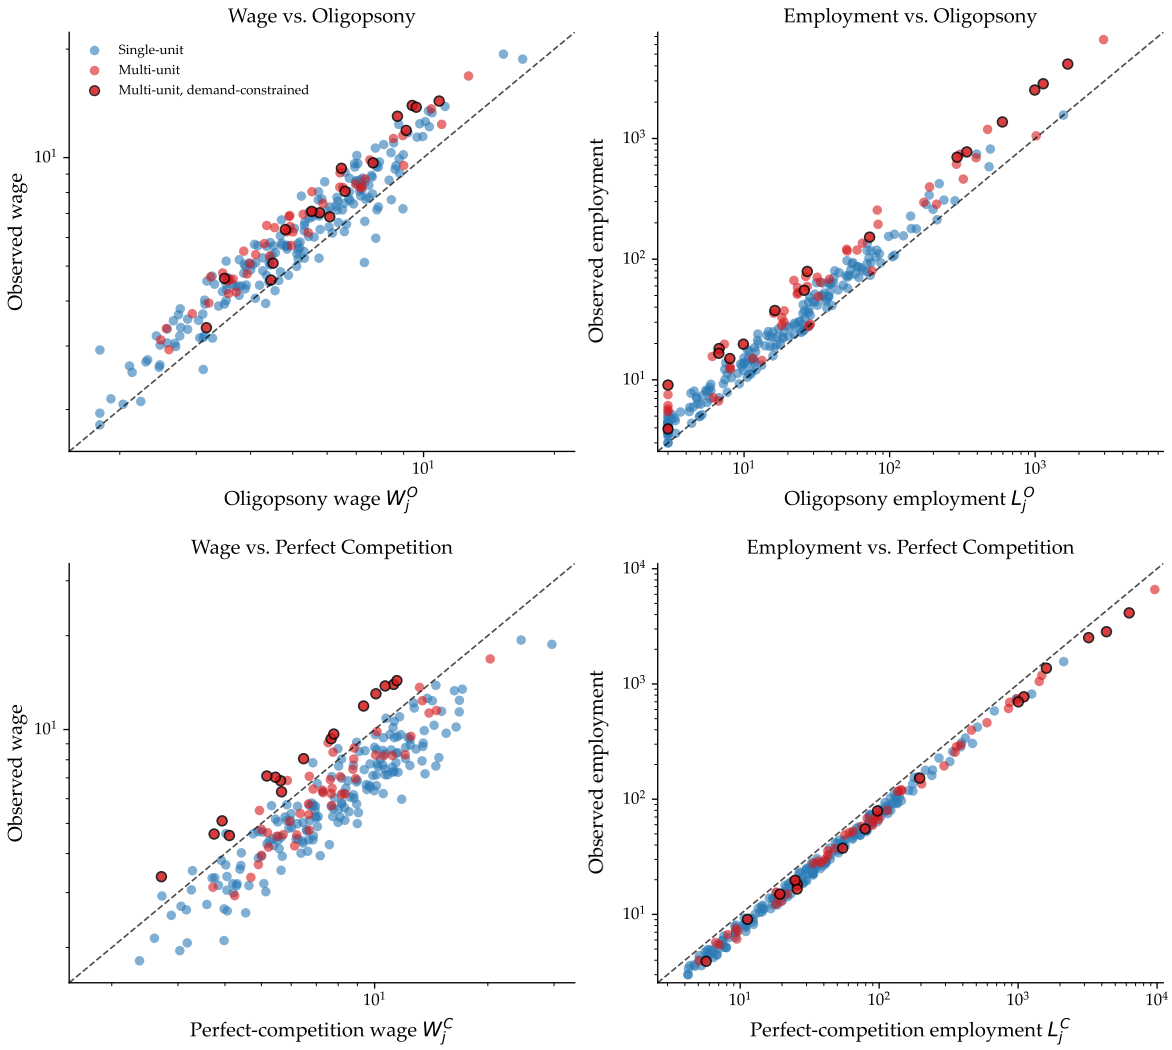


Figure 9. Distribution of Counterfactual Wage and Employment Effects, CF3 vs. Oligopsony and vs. Perfect Competition

Notes: Each union firm's observed wage (left column) and observed employment (right column) plotted (log-log) against its own oligopsony benchmark (top row) and its own perfect-competition benchmark (bottom row); the dashed line is the 45° line (no change from the benchmark). Outlined markers in the bottom row are multi-unit, demand-constrained pairs.

6 Conclusion

This paper studies how the centralization of collective bargaining shapes wages, employment, and welfare when both employers and unions have market power. We develop and estimate a model of Nash-in-Nash bargaining between firms and unions that may represent workers at several employers, in which labor supply and labor demand respond endogenously to bargained wages. Heterogeneous substitution patterns across jobs generate firm-specific monopsony power, and firm-level production functions generate firm-specific labor demand. Firms retain the right to manage, so a high enough bargained wage leads them to ration employment. In this setting, multi-unit bargaining matters through two channels: differences in bargaining power between multi- and single-unit unions, and the internalization of employment diversion across the firms a multi-unit union represents.

The main lesson for policy is that the case for centralizing bargaining depends on market conditions. Where firms are supply-constrained, as most firms in our setting are, the binding distortion is employer market power, and wages lie below the competitive level. Internalizing diversion leads a multi-unit union to moderate its wage demands, which pushes wages further below the competitive level and away from the efficient outcome. Centralization helps efficiency only if it also brings enough additional bargaining power to offset this. Where bargained wages are instead high enough to ration employment, the same moderation brings wages back toward the competitive level. Policies that promote or restrict multi-unit bargaining should therefore be evaluated against the local balance of employer and union market power, and not in isolation.

References

Ackerberg, Daniel A., Kevin Caves, and Garth Frazer, "Identification Properties of Recent Production Function Estimators," *Econometrica*, 2015, *83* (6), 2411–2451.

Angerhofer, Tirza J, Allan Collard-Wexler, and Matthew C Weinberg, "Oligopsony and Collective Bargaining," Working Paper 35476, National Bureau of Economic Research July 2026.

Arnold, David, "Mergers and acquisitions, local labor market concentration, and worker outcomes," Technical Report 2019.

Azkarate-Askasua, Miren and Miguel Zerecero, "Union and firm labor market power," Technical Report 2025.

Berger, David W., Kyle F. Herkenhoff, and Simon Mongey, "Labor Market Power," *American Economic Review*, 2022, *112* (4), 1147–1193.

Berry, Steven T., "Estimating Discrete-Choice Models of Product Differentiation," *The RAND Journal of Economics*, 1994, *25* (2), 242–262.

Bhattacharya, Vivek and Gastón Illanes, "The Design of Defined Contribution Plans," *Econometrica*, 2026. Forthcoming.

Bhuller, Manudeep, Karl Ove Moene, Magne Mogstad, and Ola L. Vestad, "Facts and Fantasies about Wage Setting and Collective Bargaining," *Journal of Economic Perspectives*, November 2022, *36* (4), 29–52.

Brancaccio, Giulia and Karam Kang, "Search Frictions and Product Design in the Municipal Bond Market," Technical Report, NYU 2025.

Card, David, "The Effect of Unions on the Structure of Wages: A Longitudinal Analysis," *Econometrica*, 1996, *64* (4), 957–979.

Collard-Wexler, Allan, Gautam Gowrisankaran, and Robin S. Lee, ""Nash-in-Nash" Bargaining: A Microfoundation for Applied Work," *Journal of Political Economy*, 2019, *127* (1), 163–195.

Crawford, Gregory S. and Ali Yurukoglu, "The Welfare Effects of Bundling in Multichannel Television Markets," *American Economic Review*, 2012, *102*, 643–685.

_ , Robin S. Lee, Michael D. Whinston, and Ali Yurukoglu, "The Welfare Effects of Vertical Integration in Multichannel Television Markets," *Econometrica*, 2018, *86*, 891–954.

Cuesta, José Ignacio, Carlos Noton, and Benjamín Vatter, "Vertical Integration Between Hospitals and Insurers," Technical Report, Stanford University 2019.

Delabastita, Vincent and Michael Rubens, "Colluding against Workers," *Journal of Political Economy*, 2025, *133* (6), 1796–1839.

DiNardo, John and David S. Lee, "Economic Impacts of New Unionization on Private Sector Employers: 1984–2001," *Quarterly Journal of Economics*, 2004, *119* (4), 1383–1441.

_ , Nicole Fortin, and Thomas Lemieux, "Labor Market Institutions and the Distribution of Wages, 1973–1992: A Semiparametric Approach," *Econometrica*, 1996, *64* (5), 1001–1044.

Dodini, Samuel, Anna Stansbury, and Alexander Willén, "How Do Firms Respond to Unions?," *Quarterly Journal of Economics*, 2026.

Dunlop, John T., *Wage Determination under Trade Unions*, New York: Macmillan, 1944.

Farber, Henry S., Daniel Herbst, Ilyana Kuziemko, and Suresh Naidu, "Unions and Inequality over the Twentieth Century: New Evidence from Survey Data," *Quarterly Journal of Economics*, 2021, *136* (3), 1325–1385.

Felix, Mayara, "Trade, Labor Market Concentration, and Wages," Working Paper 35018, National Bureau of Economic Research March 2026.

Freeman, Richard B. and James L. Medoff, "What Do Unions Do?," *Industrial and Labor Relations Review*, 1984, *38* (2), 244–263.

Galbraith, John Kenneth, "Countervailing Power," *The American Economic Review*, 1954, *44* (2), 1–6.

Gottfries, Axel and Gregor Jarosch, "Dynamic Monopsony with Large Firms and Noncompetes," Working Paper 31965, National Bureau of Economic Research 2023.

Gowrisankaran, Gautam, Aviv Nevo, and Robert Town, "Mergers When Prices Are Negotiated: Evidence from the Hospital Industry," *American Economic Review*, 2015, *105*, 172–203.

Grennan, Matthew, "Price Discrimination and Bargaining: Empirical Evidence from Medical Devices," *American Economic Review*, 2013, *103*, 145–177.

Hazell, Jonathon, Christina Patterson, Heather Sarsons, and Bledi Taska, "National Wage Setting," Working Paper 30623, National Bureau of Economic Research November 2022.

Hjort, Jonas, Xuan Li, and Heather Sarsons, "Across-Country Wage Compression in Multinationals," *American Economic Review*, January 2026, *116* (1), 52–87.

Ho, Kate and Robin S. Lee, "Insurer Competition in Health Care Markets," *Econometrica*, 2017, *85*, 379–417.

_ **and** _ , "Equilibrium Provider Networks: Bargaining and Exclusion in Health Care Markets," *American Economic Review*, 2019, *109*, 473–522.

Horn, Henrik and Asher Wolinsky, "Bilateral Monopolies and Incentives for Merger," *The RAND Journal of Economics*, 1988, *19* (3), 408–419.

Kline, Patrick M., "Labor Market Monopsony: Fundamentals and Frontiers," in "Handbook of Labor Economics," Vol. 6, Elsevier, 2025, pp. 655–728.

Kroft, Kory, Yao Luo, Magne Mogstad, and Bradley Setzler, "Imperfect Competition and Rents in Labor and Product Markets: The Case of the Construction Industry," *American Economic Review*, 2025, *115* (9), 2926–2969.

Lagos, Lorenzo, "Union bargaining power and the amenity-wage tradeoff," Technical Report, National Bureau of Economic Research 2025.

Larsen, Bradley J., "The Efficiency of Real-World Bargaining: Evidence from Wholesale Used-Auto Auctions," *Review of Economic Studies*, 2021, *88*, 851–882.

Lemieux, Thomas, "Estimating the Effects of Unions on Wage Inequality in a Panel Data Model with Comparative Advantage and Nonrandom Selection," *Journal of Labor Economics*, 1998, *16* (2), 261–291.

Lewis, H. Gregg, *Unionism and Relative Wages in the United States*, Chicago, IL: University of Chicago Press, 1963.

McDonald, Ian M. and Robert M. Solow, "Wage Bargaining and Employment," *The American Economic Review*, 1981, *71* (5), 896–908.

Milgrom, Paul and Chris Shannon, "Monotone Comparative Statics," *Econometrica*, 1994, *62* (1), 157–180.

Naidu, Suresh, Eric A Posner, and Glen Weyl, "Antitrust remedies for labor market power," *Harvard law review*, 2018, *132* (2), 536–601.

Prager, Elena and Matt Schmitt, "Employer consolidation and wages: Evidence from hospitals," *American Economic Review*, 2021, *111* (2), 397–427.

Robinson, Joan, *The Economics of Imperfect Competition*, London: Macmillan, 1933.

Robles-Garcia, Claudia, "Competition and Incentives in Mortgage Markets: The Role of Brokers," Technical Report, Stanford University 2020.

Rubens, Michael, "Management, Productivity, and Technology Choices: Evidence from US Mining Schools," *The RAND Journal of Economics*, 2023, *54* (1), 165–186.

Schubert, Gregor, Anna Stansbury, Bledi Taska, and Matthew Dey, "Employer Concentration and Wages: Evidence from Large Firms' Hiring Shocks," Technical Report 2026.

Setzler, Bradley, "Labor and Product Market Power, Endogenous Quality, and the Consolidation of the US Hospital Industry," Working Paper 34180, National Bureau of Economic Research 2025.

Sharma, Garima, "Collusion among employers in India," *Unpublished manuscript, conditionally accepted at American Economic Review*, 2025.

Appendix

A The Hiring Game: Best Responses and Uniqueness

This appendix derives the best response in equation (10) and proves Proposition 1. Fix a market and the wage vector, and write $D_j \equiv L_j^D(W_j) \geq 0$ for firm j's labor demand, with $D_j > 0$ for at least one firm. It is convenient to write the choice probabilities in equation (5) in terms of two aggregates, the inclusive value of industry g and the market-level sum:

$$x_g \equiv e^{I_g} \equiv \sum_{k \in \mathcal{J}_{g,m}} \psi_k^{1/\lambda} e^{\delta_k/\lambda}, \qquad \Sigma \equiv 1 + \sum_{g \in \mathcal{G}_m} x_g^{\lambda}, \qquad s_j = \frac{\psi_j^{1/\lambda} e^{\delta_j/\lambda}\, x_{g(j)}^{\lambda-1}}{\Sigma}. \tag{A.1}$$

Best response. Firm j's employment is $\psi_j \bar{L} s_j(\psi)$. Using equation (A.1), s_j is strictly increasing in ψ_j, so employment is a continuous, strictly increasing function of ψ_j that equals zero at $\psi_j = 0$. By the rationing condition (1), the firm hires D_j if that is attainable with $\psi_j \leq 1$, and sets $\psi_j = 1$ otherwise. When the first case applies, $\psi_j \bar{L} s_j = D_j$ gives

$$\psi_j^{\frac{1+\lambda}{\lambda}} = \frac{D_j \Sigma}{\bar{L}\, e^{\delta_j/\lambda}\, x_{g(j)}^{\lambda-1}},$$

which yields equation (10) after capping at one. The best response depends on ψ_{-j} only through $(\Sigma, x_{g(j)})$.

Equilibrium in terms of aggregates. Substituting the best responses into the definition of x_g, firm k's contribution to its industry's inclusive value is

$$t_k(x_g; \Sigma) \equiv \psi_k^{1/\lambda} e^{\delta_k/\lambda} = \min\left\{ e^{\delta_k/\lambda},\ c_k(\Sigma)\, x_g^a \right\}, \tag{A.2}$$

with

$$c_k(\Sigma) = \left(\frac{D_k \Sigma}{\bar{L}} \right)^{\frac{1}{1+\lambda}} e^{\frac{\delta_k}{1+\lambda}}, \qquad a \equiv \frac{1-\lambda}{1+\lambda} \in [0, 1),$$

where the exponent on δ_k simplifies because $\frac{1}{\lambda}\left(1-\frac{1}{1+\lambda}\right)=\frac{1}{1+\lambda}$. A vector ψ^* is an equilibrium if and only if its aggregates solve

$$x_g = h_g(x_g;\Sigma) \equiv \sum_{k\in\mathcal{J}_g} t_k(x_g;\Sigma) \quad \text{for every } g, \qquad \Sigma = 1+\sum_g x_g^\lambda, \tag{A.3}$$

since each ψ_k^* is then recovered from equation (10). It is therefore enough to show that system (A.3) has a unique solution.

Lemma 1. *For every $\Sigma \geq 1$ and every industry g with at least one firm, $x_g = h_g(x_g;\Sigma)$ has a unique positive solution $x_g(\Sigma)$.*

Proof. Each term $t_k(x;\Sigma)/x$ is the minimum of $e^{\delta_k/\lambda}/x$ and $c_k(\Sigma)x^{a-1}$. Both are strictly decreasing in x because $a<1$, so $h_g(x;\Sigma)/x$ is continuous and strictly decreasing on $(0,\infty)$. As $x\to 0$, both expressions diverge, so $h_g/x\to\infty$. As $x\to\infty$, $h_g(x;\Sigma)\leq\sum_k e^{\delta_k/\lambda}$ is bounded, so $h_g/x\to 0$. Hence $h_g(x;\Sigma)/x=1$ at exactly one $x>0$. □

Lemma 2. *For every industry, the elasticity of $x_g(\Sigma)^\lambda$ with respect to Σ lies in $[0,\frac{1}{2}]$.*

Proof. Let $\rho_g\in[0,1]$ be the share of h_g contributed by rationing firms ($\psi_k<1$) at the solution. Rationing firms' contributions scale as $\Sigma^{1/(1+\lambda)}x_g^a$, and unrationed firms' contributions do not depend on (x_g,Σ). Log-differentiating $x_g=h_g(x_g;\Sigma)$ gives

$$\frac{d\ln x_g}{d\ln\Sigma}=\frac{\rho_g/(1+\lambda)}{1-\rho_g a},$$

which is increasing in ρ_g and equals $\frac{1/(1+\lambda)}{1-a}=\frac{1}{2\lambda}$ at $\rho_g=1$. Multiplying by λ gives the bound. At the points where a firm switches regime, h_g has a kink and the same bound holds for the one-sided derivatives. □

Proof of Proposition 1. By Lemma 1, the industry conditions in system (A.3) define $x_g(\Sigma)$ for every Σ, and the market condition becomes $\Sigma = G(\Sigma)\equiv 1+\sum_g x_g(\Sigma)^\lambda$. By Lemma 2, the elasticity of G with respect to Σ is at most $\frac{1}{2}\cdot\frac{G-1}{G}<\frac{1}{2}$, so $G(\Sigma)/\Sigma$ is strictly decreasing. Since some firm has positive demand, $x_g(1)>0$ for some g and $G(1)/1>1$. Since G grows at most like $\sqrt{\Sigma}$, $G(\Sigma)/\Sigma\to 0$ as $\Sigma\to\infty$. Hence there is exactly one Σ^* with $G(\Sigma^*)=\Sigma^*$, the inclusive values $x_g(\Sigma^*)$ are unique, and so is the equilibrium vector ψ^* given by equation (10). □

Computation. The proof is also the solution algorithm. For a trial Σ, we solve the scalar equation $x_g = h_g(x_g; \Sigma)$ for each industry by bracketing, and then solve the scalar equation $\Sigma = G(\Sigma)$ in an outer loop. Both functions are monotone in the relevant sense, so standard bracketing methods converge. The minimum in equation (A.2) classifies firms into regimes along the way. Two special cases have closed forms and serve as checks. If every firm in the market rations, then $x_g^\lambda = \Omega_g \sqrt{\Sigma}$ with $\Omega_g \equiv \bar{L}^{-1/2} \left[\sum_{k \in \mathcal{J}_g} (D_k e^{\delta_k})^{1/(1+\lambda)}\right]^{(1+\lambda)/2}$, and the market condition is a quadratic in $\sqrt{\Sigma}$ with unique positive root $\sqrt{\Sigma^*} = \left(\Omega + \sqrt{\Omega^2 + 4}\right)/2$, where $\Omega = \sum_g \Omega_g$. When $\lambda = 1$, the model collapses to a plain logit without nests, treated next.

The case without nests ($\lambda = 1$). Fix a market and let $\delta_j = X_j'\phi + \beta W_j$ denote mean utility net of any rationing (observable job characteristics and wage only). For a vector of rationing shares $\psi = (\psi_1, \ldots, \psi_J)$, define the aggregator

$$\Sigma(\psi) = 1 + \sum_k \psi_k \exp(\delta_k),$$

so that the aggregate labor supply facing firm j – accounting for the fact that only a share ψ_j of interested workers is actually hired – is

$$S_j(\psi) = \bar{L} \cdot \frac{\psi_j \exp(\delta_j)}{\Sigma(\psi)}.$$

Market clearing requires that realized employment at firm j, $\psi_j S_j(\psi)$, never exceed labor demand $D_j = L_j^D(W_j)$ from Section 4.2:

$$\psi_j = \min\left\{1,\ \frac{D_j(W_j)}{S_j(\psi)}\right\}.$$

A firm with $\psi_j = 1$ is unrationed – the usual supply-constrained (oligopsony) case – while a firm with $\psi_j < 1$ is rationing employment down to its labor demand, the demand-constrained region of Figure 1.

Substituting the market-clearing condition into $S_j(\psi)$ gives a closed-form best response for every firm as a function of the aggregator alone,

$$\psi_j^*(\Sigma) = \min\left\{1,\ \sqrt{\frac{D_j(W_j)\,\Sigma}{\bar{L}\exp(\delta_j)}}\right\},$$

so the entire system of J rationing shares collapses to a single scalar fixed point, $\Phi(\Sigma) \equiv 1 + \sum_k \psi_k^*(\Sigma)\exp(\delta_k) - \Sigma = 0$. Since Φ is strictly decreasing in Σ, this fixed point is unique, and the equilibrium rationing shares ψ^* follow in closed form from Σ^*. We therefore solve for Σ^* with one-dimensional root-finding rather than a joint fixed-point iteration over all J rationing shares.

B Labor Supply Estimation Details

Utility, Shares, and Supply Elasticities

This subsection gives the labor supply elasticities implied by the nested logit specification in Section 4.1. All derivatives are taken holding the rationing shares ψ fixed; for supply-constrained firms ($\psi_j = 1$), the relevant case for the markdown, they are the elasticities of L_j^S in equation (6).

Own-wage labor supply elasticity. Differentiating equation (5) with respect to W_j,

$$e_j \equiv \frac{\partial L_j^S}{\partial W_j} \cdot \frac{W_j}{L_j^S} = \beta W_j \left[\frac{1 - s_{j|g(j)}}{\lambda} + s_{j|g(j)}\left(1 - s_{g(j)}\right) \right]. \tag{B.1}$$

This elasticity is always positive (upward-sloping labor supply) and increasing in the wage sensitivity β. It is larger the more closely workers substitute within industries (lower λ) and the smaller the firm's share of its industry, which gives scope for firms with large within-industry shares to mark down wages.

Cross-wage labor supply elasticities. The cross-wage elasticities are

$$e_{kj}^{(\text{same nest})} = \beta W_j \, s_{j|g(j)} \left[1 - s_{g(j)} - \tfrac{1}{\lambda} \right], \quad g(k) = g(j),\ k \neq j, \tag{B.2}$$

$$e_{kj}^{(\text{diff. nest})} = -\beta W_j \, s_j, \quad g(k) \neq g(j). \tag{B.3}$$

Both are negative, since jobs are substitutes. The same-nest elasticity is larger in magnitude than the different-nest one whenever $\lambda < 1$, so employment diversion, and with it the multi-unit term in equation (2), is concentrated among firms in the same industry.

Estimation Algorithm

This subsection gives the estimation algorithm summarized in Section 4.1.1. Because the rationing shares ψ depend on $\theta = (\phi, \beta, \lambda)$ through the hiring game, and the left-hand side of equation (7) depends on ψ, we iterate between the two:

1. Start from $\psi_j = 1$ for every firm and an initial θ.
2. Given ψ, recover choice probabilities from observed employment, $s_j = L_j/(\psi_j \bar{L})$, and compute the within-industry shares $s_{j|g(j)}$ and the outside share s_{0m}.
3. Estimate θ by two-stage least squares on equation (7), instrumenting W_j and $\ln s_{j|g(j)}$ with the instruments in Section 4.1.1.
4. Given the new θ, compute $\delta_j = X_j'\phi + \beta W_j$ and solve the hiring game (Appendix A) in each market for the new ψ. Non-union firms set the oligopsony wage and are supply-constrained by construction, so $\psi_j = 1$ for them.
5. Update (θ, ψ) with damping and return to step 2 until both converge.

C Production Function Estimation Details

This appendix gives the full Ackerberg et al. (2015) (ACF) estimation algorithm summarized in Section 4.2.

Given the Leontief structure of the production function in materials, a necessary condition for the firm to be profit-maximizing is that it always chooses materials such that the Leontief constraint binds, i.e., $R_j = F(K_j, L_j)\exp(\omega_j) = \beta^M M_j$. Applying logarithms, denoting the logarithm of a variable with a lowercase letter, and using a translog for $F(\cdot)$, we can write the production function as:

$$r_j = f(k_j, l_j) + \omega_j + \varepsilon_j\,, \qquad f(k_j, l_j) = \beta_l l_j + \beta_k k_j + \beta_{ll} l_j^2 + \beta_{kk} k_j^2 + \beta_{lk} l_j k_j\,,$$

where we also assume an AR(1) process for the productivity shock,

$$\omega_j = \rho\omega_{j-1} + \eta_j, \quad \mathbb{E}[\eta_j] = 0\,.$$

OLS estimation of this equation is biased because labor (and materials) respond to the firm's private information about ω_j, which the econometrician does not observe. We address this simultaneity problem using the control function approach of Ackerberg et al. (2015).

Identification rests on three assumptions. First, *invertibility*: material demand is strictly monotonic in ω_j conditional on (k_j, l_j), so that observed materials use can be inverted into a nonparametric proxy for the unobserved productivity shock, purging r_j of ω_j in a first stage without yet identifying the coefficients of $f(\cdot)$. Second, *timing*, already stated in Section 4.2: capital is fixed one period in advance of ω_j, while labor is chosen after ω_j is realized but before the revenue shock ε_j; this is what lets materials (also chosen after ω_j) proxy for productivity without also absorbing variation in labor, and it is the same timing we impose on the bargaining model, so the labor input recovered here is the one relevant for the marginal-revenue-product objects entering the firm's bargaining problem. Third, *Markovian productivity*, via the AR(1) law of motion above: the labor terms of $f(\cdot)$ are identified in a second stage from the correlation between labor and the productivity innovation η_j, net of the component of productivity predictable from lagged capital and labor, using past (lagged) employment and materials as instruments – valid as long as η_j is realized after the previous period's labor and capital decisions, so that it is uncorrelated with these lagged choices. Together these assumptions let us recover $f(\cdot)$, and with it the firm-level elasticities β_j^L and ε_j^D of Section 4.2, without an explicit instrument for current labor, at the cost of the timing and Markovian restrictions above. Table A1 reports the estimated coefficients.

D Recovering Bargaining Power

This appendix derives the first-order conditions of the Nash-in-Nash problem in equation (12) in each regime of the firm, and the closed-form recovery of bargaining power summarized in Section 5.1. Write $S_{uj}(W_j)$ for the union's gain from trade in equation (12), and $\phi_j^u \equiv d \ln S_{uj} / d \ln W_j$ for its elasticity at the agreement, given in equation (13). At the agreement $S_{uj} = W_j L_j$, since the change in the wage bill at the union's other firms is zero there. With $\bar{\pi}_j = 0$, the log of the Nash product is $(1 - \gamma_j^u) \ln \pi_j + \gamma_j^u \ln S_{uj}$, and a bargained wage at which the objective is differentiable satisfies

$$\gamma_j^u \, \phi_j^u = \left(1 - \gamma_j^u\right) \chi_j, \qquad \chi_j \equiv -\frac{d \ln \pi_j}{d \ln W_j}, \tag{D.1}$$

Table A1. Translog Production Function Coefficients

	ACF, Translog
$\hat{\beta}_l$	1.2257^{***}
	(0.0000)
$\hat{\beta}_k$	-0.2691^{***}
	(0.0000)
$\hat{\beta}_{ll}$	-0.0484^{***}
	(0.0000)
$\hat{\beta}_{kk}$	0.0243^{***}
	(0.0000)
$\hat{\beta}_{lk}$	-0.0331^{***}
	(0.0000)

Notes: Coefficients of the value-added translog $f(k_j, l_j)$, estimated following Ackerberg et al. (2015) with materials as the productivity proxy. Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. The implied firm-level output and labor demand elasticities are summarized in Table 4.

where χ_j is the elasticity of the firm's profit loss with respect to the wage. The two regimes of equation (1) differ in how employment, and therefore profit, responds to the wage.

Supply-constrained regime ($\psi_j = 1$). Employment is $L_j = L_j^S(W_j)$, with elasticity e_j, so

$$\frac{d\pi_j}{dW_j} = \left(R_j'(L_j) - W_j\right)\frac{dL_j^S}{dW_j} - L_j = L_j\left[\left(\frac{R_j'(L_j)}{W_j} - 1\right)e_j - 1\right].$$

Write $W_j = \lambda_j R_j'(L_j)$. Since $R_j'(L_j) = \beta_j^L R_j / L_j$, profit is $\pi_j = R_j'(L_j)L_j\left(1/\beta_j^L - \lambda_j\right)$, and

$$\chi_j = \frac{\lambda_j(1+e_j) - e_j}{1/\beta_j^L - \lambda_j}.$$

Substituting into equation (D.1) and solving for λ_j gives

$$\lambda_j = \frac{\gamma_j^u \phi_j^u / \beta_j^L + (1-\gamma_j^u)\, e_j}{\gamma_j^u \phi_j^u + (1-\gamma_j^u)(1+e_j)} = \omega_j^u \frac{1}{\beta_j^L} + \left(1-\omega_j^u\right)\mu_j,$$

which is equation (3), with ω_j^u and $\mu_j = e_j/(1+e_j)$ as defined there. Given the observed wedge $W_j / R_j'(L_j)$, this expression is linear in ω_j^u and delivers equation (16). Under the wage bill objective, $\phi_j^u = (1+e_j)(1+\kappa_j)$, so ω_j^u takes the form in equation (15), which is strictly increasing in γ_j^u for $\kappa_j > -1$ and inverts to equation (17).

Demand-constrained regime ($\psi_j < 1$). Employment is $L_j = L_j^D(W_j)$ and the wage equals the marginal revenue product, $W_j = R_j'(L_j)$. Since the firm is on its labor demand curve, the envelope theorem gives $d\pi_j/dW_j = -L_j$. Moreover, $W_j L_j = \beta_j^L R_j$ and $\pi_j = (1-\beta_j^L)R_j$, so $\chi_j = W_j L_j/\pi_j = \beta_j^L/(1-\beta_j^L) \equiv b_j$, and equation (D.1) becomes

$$\gamma_j^u \, \phi_j^u = \left(1-\gamma_j^u\right) b_j, \tag{D.2}$$

with ϕ_j^u evaluated at $\epsilon^{L_j,W_j} = \varepsilon_j^D$. Solving for γ_j^u gives equation (18). Since $b_j > 0$, equation (D.2) has a solution only if $\phi_j^u > 0$: the union must still gain from a higher wage even though employment falls. For a single-unit union, $\phi_j^u = 1 + \varepsilon_j^D$, so this requires labor demand to be inelastic at the bargained wage, $\varepsilon_j^D > -1$. If labor demand is more elastic, raising the wage above the competitive level lowers the wage bill, and no amount of bargaining power pushes the firm into rationing.

The competitive wage. At the competitive wage W_j^C, labor supply equals labor demand, $\psi_j = 1$, and $\lambda_j = 1$. The Nash objective has a kink at W_j^C, because employment follows labor supply to the left and labor demand to the right. Evaluating the supply-constrained χ_j at $\lambda_j = 1$ gives $\chi_j = b_j$, the same as in the demand-constrained regime, so the one-sided derivatives differ only through the union's elasticity: $\phi_j^{u,S}$, evaluated at e_j, to the left, and $\phi_j^{u,D}$, evaluated at ε_j^D, to the right. The competitive wage is the bargained wage when the objective is increasing to the left and decreasing to the right, $\gamma_j^u \phi_j^{u,S} \geq (1-\gamma_j^u) b_j \geq \gamma_j^u \phi_j^{u,D}$, that is,

$$\frac{b_j}{\phi_j^{u,S} + b_j} \;\leq\; \gamma_j^u \;\leq\; \frac{b_j}{\phi_j^{u,D} + b_j}, \tag{D.3}$$

where the upper bound is one if $\phi_j^{u,D} \leq 0$. For a single-unit union, $\phi_j^{u,S} = 1 + e_j$ and $\phi_j^{u,D} = 1 + \varepsilon_j^D$, and the interval becomes

$$\frac{\beta_j^L}{1 + (1-\beta_j^L)\, e_j} \;\leq\; \gamma_j^u \;\leq\; \frac{\beta_j^L}{1 + (1-\beta_j^L)\, \varepsilon_j^D},$$

with both elasticities evaluated at the competitive employment level. Below the lower bound the wage is supply-constrained and given by equation (3); above the upper bound the firm rations and the wage solves equation (D.2). For pairs whose observed wage equals the competitive wage,

bargaining power is set-identified on the interval in equation (D.3).