Conduct in U.S. Labor Markets*

Michael Rubens[†] Bradley Setzler[‡] Chen Yeh[§]

This draft: August, 2026. Preliminary; please do not distribute.

Abstract

Over the past decade, a broad consensus has emerged that labor markets are not perfectly competitive. The natural question is: what assumption replaces perfect competition in labor markets? We develop a scalable framework for testing which labor market conduct best fits the data. Our approach distinguishes among perfect competition, monopsonistic competition, oligopsonistic competition, and collusion, as well as partially-collusive conducts that combine oligopsonistic and collusive elements. Our approach is scalable in the sense that it can be applied on a case-by-case basis to each local labor market observed in administrative data using only five commonly-available variables: employment, wages, revenues, a market identifier, and a labor demand shifter. Applying the framework to each of 108,000 labor markets defined by industry-commuting zone pairs in U.S. Census micro-data, we find that 35 percent of employment is in markets best characterized by monopsonistic competition, 40 percent by oligopsonistic competition, 22 percent by full collusion, and 3 percent by partial collusion. By allowing each market to operate under its best-fitting conduct, we find nationally that markdowns are stronger on average and more dispersed than any single conduct assumption would imply. Collusive conduct is concentrated in markets with few competitors and more prevalent in markets with lower entry rates and less differentiated employers. Applying our conduct test to seven specific markets investigated by the U.S. Department of Justice for employer collusion, we detect collusion in five, including both cases that led to a guilty plea or conviction.

---

*Any views expressed are those of the authors and not those of the U.S. Census Bureau or the Federal Reserve Bank of Richmond. The Census Bureau has reviewed this data product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data used to produce this product. This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 2567. (CBDRB-FY26-025; CBDRB-FY26-P2567-R13113). Any references to collusion in this research are strictly in the economic and not legal sense. There are no relevant sources of funding, financial relationships, or conflicts of interest to disclose. For helpful comments and suggestions, we thank Matt Weinberg (discussant), John Asker, Sydnee Caldwell, Alan Manning, Alex Mas, Joris Pinkse, Chris Sullivan, and participants at Oslo University, Tilburg University, the Oslo Labor Workshop, the IO Society meetings at ASSA, the CCE Workshop on Competition Policy, the IAB Workshop on Imperfect Labor Market Competition, and SITE Micro and Macro of Labor Markets.

[†]Boston College, NBER, and CEPR.

[‡]Ohio State University and NBER.

[§]Federal Reserve Bank of Richmond.

1 Introduction

A broad consensus has emerged that labor markets are not perfectly competitive. Labor supply is not perfectly elastic across jobs and firms set wages below the marginal revenue product of labor (MRPL). Recent surveys by Card (2022), Azar and Marinescu (2024), Kline (2025), and Caldwell, Dube, and Naidu (2026) document this consensus and the accumulating evidence behind it. The natural question is: *what assumption replaces perfect competition in labor markets?* The recent literature has proposed several alternatives. One strand assumes monopsonistic competition, in which firms set wages in consideration of the firm-specific labor supply elasticity, ignoring strategic interactions across competitors (Card, Cardoso, Heining, and Kline, 2018; Lamadon, Mogstad, and Setzler, 2022). A second strand assumes oligopsonistic competition, in which firms anticipate the strategic responses of rivals when making wage and employment decisions (Berger, Herkenhoff, and Mongey, 2022; Azar, Berry, and Marinescu, 2022). A third strand allows for collusion among employers to suppress wages (Delabastita and Rubens, 2025; Sharma, 2025). Each of these alternatives constitutes a different conduct. With few exceptions, the literature has taken a particular conduct as given, rather than using data to learn the conduct.[1]

The choice of conduct is consequential for both measurement and policy. Different conduct assumptions imply different levels and dispersion of wage markdowns—the key measure of labor market power in the literature. For example, oligopsonistic conducts allow firms with greater market share to impose stronger markdowns, while monopsonistic competition does not, as emphasized by Berger et al. (2022) and a subsequent literature. These differences propagate into aggregate calculations of misallocation and welfare losses due to labor market imperfections. The choice of conduct also has direct implications for antitrust policy. Under oligopsonistic competition, the relevant policy tools might be merger review and structural remedies aimed at market concentration. Under collusion, the relevant tools might be prosecution of wage-fixing and no-poach agreements. Using the wrong conduct assumption can therefore lead to misleading welfare calculations and misdirected policy.

In this paper, we develop a practical, data-driven framework for determining which conduct best characterizes each local labor market and apply it to administrative data from the United States. A central design constraint we impose on ourselves is scalability: the approach must be feasible for tens of thousands of markets—including in service industries, where much of the economy operates—using only five variables commonly available in tax administration and census datasets. These variables are employment, wages, revenues, a market identifier, and a labor demand shifter. The core empirical strategy proceeds by constructing markdowns using both the supply-side of the

[1] Exceptions include Roussille and Scuderi (2026), who test among perfect, monopsonistic, and oligopsonistic competition in an online tech labor market, and Delabastita and Rubens (2025), who test between full collusion and unilateral Cournot oligopsony in a 19th-century coal-mining cartel.

problem—that the markdown is optimally a conduct-specific function of labor supply parameters—and the demand-side of the problem—that the markdown can be expressed as the ratio of the wage to the MRPL, regardless of conduct. We construct a test statistic that compares how the conduct-free demand-side markdown and the conduct-dependent supply-side markdown respond to rivals' shocks, where these shocks serve as instrumental variables (IV). Intuitively, an assumed conduct fits the data well if its conduct-dependent supply-side markdowns covary with the IV to the same extent as the conduct-free demand-side markdowns across environments.

The model combines two standard ingredients from the labor economics literature. On the supply side, workers have horizontally-differentiated preferences over jobs, generating upward-sloping firm-level residual labor supply curves. On the demand side, firms produce revenue using a technology that exhibits a diminishing marginal revenue product of labor. The combination of supply and demand primitives, together with a conduct assumption, allows us to construct both supply-side and demand-side markdowns, up to specification and measurement errors. We consider eight alternative conducts, arising from the combination of the mode of competition—perfect competition, monopsonistic competition, Bertrand oligopsony, and Cournot oligopsony—with the possibility of collusion—either partially or fully, in both Bertrand and Cournot variants. Figure 1 visualizes this taxonomy of conducts. To our knowledge, this taxonomy spans the union of conducts that have been considered in the recent static wage-posting literature, while also introducing several that have not been considered. For example, we are not aware of any prior work on partially-collusive labor market conducts, combining oligopsonistic and collusive elements. Our taxonomy also spans the union of conducts considered in leading IO articles on conduct-testing in product markets (e.g., Miller and Weinberg 2017, Backus, Conlon, and Sinkinson 2021, Duarte, Magnolfi, Sølvsten, and Sullivan 2024).

To discriminate among the eight candidate conducts, we adapt the Rivers and Vuong (2002) testing procedure to the labor market setting. Each candidate conduct implies a different set of wage markdowns for every firm in a given market. These markdowns enter a moment condition that captures covariation in supply-side and demand-side markdowns in response to rivals' shocks. Pairwise comparison of the GMM objective function across candidate conducts identifies which provides the best fit, and the Rivers-Vuong test statistic provides a formal basis for determining whether the difference is statistically significant. We apply this procedure separately to each of approximately 108,000 local labor markets—defined as 3-digit NAICS industry by commuting zone pairs, following Berger et al. (2022)—using establishment-level micro-data from the U.S. Census Bureau. Establishment-level revenues are drawn from the quinquennial Economic Censuses from 1987 to 2017, covering industries that account for roughly 70 percent of private non-farm employment. Employment and payroll come from the Longitudinal Business Database, which provides annual data on the near-universe of private-sector employer establishments. Labor demand

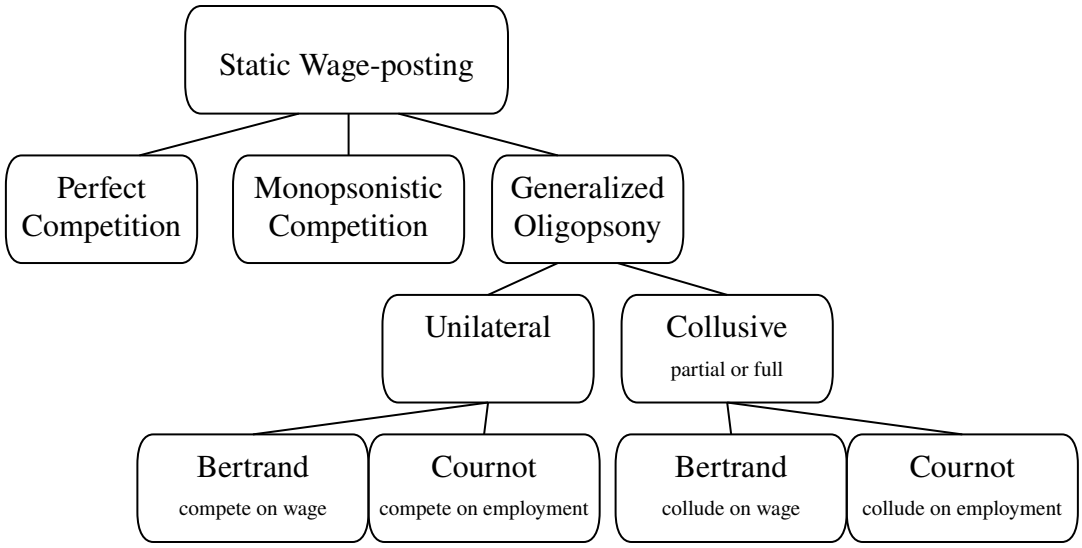


Figure 1: Taxonomy of Static Wage-posting Labor Market Conducts

Notes: This tree classifies the eight candidate conduct assumptions tested in the paper. At the top level, we restrict attention to static models of wage-posting. The second level distinguishes conducts without strategic interactions among competitors (perfect and monopsonistic competition) from those with strategic interactions (generalized oligopsony). The third level separates unilateral oligopsony from collusive oligopsony, where the degree of collusion may be partial or full. The fourth level distinguishes between Bertrand competition (competing or colluding on wages) from Cournot competition (competing or colluding on employment).

shocks—the IV we use to identify labor supply parameters—are constructed from regional housing price variation following Giroud and Mueller (2019), exploiting cross-location propagation through firms' multi-establishment networks.

Across all U.S. markets, weighted by employment, we find that 35 percent are best characterized by monopsonistic competition, 40 percent by oligopsonistic competition, 22 percent by full collusion, and 3 percent by partial collusion. When restricting to markets in which our best-fit conduct performs significantly better (in a statistical sense) than the other conduct types, we find evidence for full and partial collusion for 18 and 1 percent of employment, respectively. Allowing for this heterogeneous conduct across markets generates a markdown distribution that is both lower on average and more dispersed than the distribution implied by any single assumed conduct applied uniformly. This result resolves a key puzzle in the literature: demand-side markdowns, like those estimated by Yeh, Macaluso, and Hershbein (2022) and Rubens (2023), tend to be lower and more dispersed than supply-side markdowns that are based on assuming a national unilateral conduct. Incorporating conduct heterogeneity in local labor markets brings the national supply-side markdown distribution in line with the national demand-side markdown distribution.

In non-collusive markets, we find that much of the variation in markdowns is driven by heterogeneity in the labor supply elasticity in the market rather than by cross-firm heterogeneity in market

shares. This is because only 1 percent of employers hold a market share of 10 percent or more, the range where oligopsony allows for large firms to set substantially stronger markdowns. As external validation, we apply the conduct test to seven specific markets that have been investigated by the U.S. Department of Justice (DOJ) for criminal employer collusion. Our test detects collusion in five of the seven cases, including both that led to a conviction or guilty plea, compared to a baseline detection rate of 22 percent across all markets.

Our paper relates to three strands of the literature. The first is a literature which has developed and estimated models of imperfect competition in labor markets in which markdowns can be inferred from labor supply parameters under a maintained conduct assumption. Within this literature, the monopsonistic competition framework has been applied by Card et al. (2018), Lamadon et al. (2022), and Kroft, Luo, Mogstad, and Setzler (2025), while the oligopsonistic competition framework has been developed by Berger et al. (2022), Azar et al. (2022), Rubens (2023), Chan, Kroft, Mattana, and Mourifié (2024), and Deb, Eeckhout, Patel, and Warren (2024). Our contribution relative to this literature is to test between these different conduct assumptions using a rich set of alternatives, rather than maintaining any one of them. We follow most of the recent literature on imperfectly competitive labor markets in restricting attention to static wage-posting conducts, which Manning (2021) calls "New Classical," and do not consider dynamic oligopsony models, such as that of Jarosch, Nimczik, and Sorkin (2024), or union bargaining models, such as that of Angerhofer, Collard-Wexler, and Weinberg (2025).[2]

The second related strand is the conduct-testing literature, which originated in industrial organization (IO) with the work of Bresnahan (1982) and Berry and Haile (2014), and has been applied in product market settings by Miller and Weinberg (2017), Backus et al. (2021), and Duarte et al. (2024). Recent work by Roussille and Scuderi (2026) and Delabastita and Rubens (2025) has extended conduct testing to case studies of specific labor markets. Our contribution is to test among conduct types at scale, across thousands of U.S. labor markets in population administrative data, as well as to allow for decreasing MRPL and consider a broader menu of alternative conducts. Our introduction of diminishing MRPL to labor-market conduct-testing is analogous, though it differs in some important ways, to the recent introduction of increasing marginal cost to the product-market conduct-testing literature by Duarte, Magnolfi, Quint, Sølvsten, and Sullivan (2025).

The third related strand concerns the detection and antitrust implications of oligopsony and collusion. This literature has gained particular importance since the U.S. Department of Justice (DOJ) and Federal Trade Commission (FTC) incorporated anticompetitive labor market activity into the U.S. Merger Guidelines (US DOJ and FTC, 2023). The empirical analysis of mergers under oligop-

[2]Unionized labor markets are rare in the US: only 6% of private-sector jobs are unionized (BLS, 2025). Lachowska, Mas, Saggio, and Woodbury (2022) find that employer-employee bargaining over wages is rare in the US. In Germany, Caldwell, Haegele, and Heining (2026) find that bargaining is common among specialized, experienced, high-wage workers, but the wage margins over which they bargain are small in magnitude.

sony has been advanced by Prager and Schmitt (2021), Berger, Hasenzagl, Herkenhoff, Mongey, and Posner (2025), Arnold (2025), and Setzler (2025), while the literature on employer collusion and labor antitrust includes Masur and Posner (2023), Gibson (2025), Gottfries and Jarosch (2025), Sharma (2025), and Herrera-Caicedo, Jeffers, and Prager (2026). Our first contribution relative to this literature is to apply the conduct test to specific markets flagged by the DOJ, demonstrating that our approach detects collusive behavior in many of the same markets that the DOJ accused of collusion. Thus, our approach provides an empirical bridge between the economic analysis of employer coordination and the legal framework for labor antitrust. Our second contribution concerns using market-level characteristics to predict collusion. A classical theoretical literature in IO argues that particular characteristics of markets—especially a small number of competitors, a low entry rate, and little (product) differentiation—are predictive of collusive behavior (Bain, 1956; Stigler, 1964; Scherer, 1980; Deneckere, 1983; Rotemberg and Saloner, 1986; Ross, 1992). However, our study appears to be the first to take these predictions to the data at scale, in labor or product contexts. We do so by first estimating the conduct of many markets and then correlating the presence of collusion with these market characteristics, confirming that collusion is predicted by a small number of competitors, a low entry rate, and little (employer) differentiation

The remainder of the paper is organized as follows. Section 2 presents the model. Section 3 develops identification and testing. Section 4 describes the empirical specification. Section 5 describes the data sources and implementation details. Section 6 presents the main results. Section 7 explores the antitrust implications. Section 8 concludes.

2 Model

This section presents the model primitives, the taxonomy of labor market conducts, and the mapping from conduct to wage markdowns.

2.1 Worker Preferences and Labor Supply

Consider an economy with workers indexed by i, firms indexed by j or f, time periods indexed by t, and local labor markets indexed by m. A market is defined by a set of firms $\mathcal{J}_{mt}$ and a set of

workers $\mathcal{I}_{mt}$.[3] Let W_{jt} and L_{jt} denote the wage and employment of firm j at time t, respectively.[4]

Each worker's problem is to choose the firm that provides the greatest utility:

$$j^*(i,t) = \arg\max_{j \in \mathcal{J}_{m(i)t}} U_{ijt}, \quad \text{where} \quad U_{ijt} = \mathcal{U}(W_{jt}, \xi_{jt}, \eta_{ijt}; \boldsymbol{\alpha}). \tag{1}$$

Here, ξ_{jt} represents vertical differentiation across firms that is unobserved by the econometrician, which we refer to as amenities. The term η_{ijt} represents horizontal differentiation that is unobserved by both the firm and the econometrician, which we refer to as match-specific tastes. Unlike vertical amenities—where workers uniformly prefer more to less—these tastes capture idiosyncratic preferences over which workers naturally disagree. Canonical examples include the worker-job commuting distance, the alignment between a worker's need for schedule flexibility and the job's specific hours, or a worker's subjective fit with the firm's corporate culture and managerial style. The parameter vector $\boldsymbol{\alpha}$ governs workers' sensitivity to wages.

The worker's discrete choice problem, combined with the distribution of match tastes, generates an upward-sloping residual labor supply curve for each firm. The market share of firm j is the probability that a worker in market m chooses firm j, integrated over the distribution of match-specific tastes:

$$s_{jt} = \int \mathbf{1}\left[\mathcal{U}(W_{jt}, \xi_{jt}, \eta_{ijt}; \boldsymbol{\alpha}) \geq \mathcal{U}(W_{ft}, \xi_{ft}, \eta_{ift}; \boldsymbol{\alpha}) \;\; \forall f \in \mathcal{J}_{m(j)t} \cup \{0\}\right] dF(\boldsymbol{\eta}_i). \tag{2}$$

We write this more compactly as $s_{jt} = S_{jt}\left(\{W_{ft}\}_{f \in \mathcal{J}_{m(j)t}}\right)$. We define the inside market share (which is observed) as $s_{jt|m} \equiv L_{jt}/(\sum_{f \in \mathcal{J}_{m(j)t}} L_{ft})$, the outside share (which is unobserved) as $s_{0t|m} \equiv 1 - \sum_{f \in \mathcal{J}_{mt}} s_{ft}$, and the total market share (which is unobserved) as $s_{jt} \equiv s_{jt|m}(1 - s_{0t|m})$. Raising the wage attracts workers from competing firms as well as from the outside option, so if α is finite, each firm faces an upward-sloping labor supply curve.

[3] In theory, markets can be defined arbitrarily. For example, markets could be based on geography, industry, occupation, firm size, worker experience, or any other characteristics of workers or firms. In our empirical application, we define markets in terms of geography and industry, following Berger et al. (2022), Lamadon et al. (2022), Yeh et al. (2022), and a subsequent literature. Our application considers establishments rather than firms, to account for the fact that some firms operate establishments in multiple locations. We follow the convention of using the term "firm" rather than "establishment" in the theory section of this paper with slight abuse of terminology.

[4] In the main text, we assume that the researcher only has access to data on total employment and wages at the firm-level, consistent with the empirical context considered by, e.g., Berger et al. (2022) and Yeh et al. (2022). However, as we show in Appendix E, our model and testing framework can easily be extended to account for worker heterogeneity if data on worker types are available. For example, it is straightforward to allow that competitive conduct is different for college-educated and non-college-educated workers within the same market, given worker-level education data.

2.2 Revenue Production and Profits

Each firm generates revenue $R_{jt} = F_{jt}(L_{jt}, M_{jt}, K_{jt})$, which depends on labor L_{jt}, materials M_{jt}, and capital K_{jt}. The marginal revenue product of labor is defined as

$$MRPL_{jt}(L_{jt}, M_{jt}, K_{jt}) \equiv \frac{\partial F_{jt}(L_{jt}, M_{jt}, K_{jt})}{\partial L_{jt}} > 0. \tag{3}$$

We allow for diminishing marginal revenue product: $\frac{\partial}{\partial L_{jt}} MRPL_{jt}(L_{jt}, M_{jt}, K_{jt}) \leq 0$. This is a standard assumption in the labor economics literature, and it is analogous to allowing for increasing marginal cost curves in the industrial organization conduct-testing literature (Duarte et al., 2025).

Firm profits Π_{jt} are defined as the difference between revenues and input expenditures:

$$\Pi_{jt} \equiv R_{jt} - L_{jt}W_{jt} - \rho M_{jt} - rK_{jt}. \tag{4}$$

We follow the literature on production function estimation in assuming that materials are chosen simultaneously with employment, while capital is chosen in the previous year, and that both are purchased in competitive markets. A data limitation that we impose is that capital and materials are not observed by the econometrician. Later, we present a practical estimation strategy that allows us to account for endogenous capital and materials despite not observing either variable, thus overcoming a common data limitation for service industries.

2.3 The Firms' Problem under Alternative Conducts

The labor market conduct determines exactly how firms set wages and employment to maximize the profits defined above. We consider eight candidate conducts, arising from the interaction between the mode of competition and the degree of employer collusion.

Table 1 summarizes the firm's optimization problem under each assumed conduct. The first two conducts assume that there are no strategic interactions in the labor market. This means that firms assume that their competitors will not respond to their choices of wages and employment. Under perfect competition, the firm takes its wage as a given constant and chooses its employment only. Under monopsonistic competition, the firm chooses its wage-employment pair along the given labor supply curve.

The next two conducts assume that firms are unilaterally oligopsonistic. Under Cournot oligopsony, the firm chooses employment taking rivals' employment as fixed. Since rivals' employment levels do not change, this is equivalent to the reference firm assuming that rivals' wages adjust by exactly the amount needed to retain their workers, so that attempted expansion by the reference firm fails to poach workers from competitors. Thus, the reference firm assumes all of its new

Table 1: Firm's Problem under Each Conduct

Conduct	Firm's Problem
Perfect competition	$\max_{L_{jt}} \Pi_{jt}\big\vert_{W_{jt},\{L_{-jt},W_{-jt}\}}$
Monopsonistic competition	$\max_{W_{jt}} \Pi_{jt}\big\vert_{\{L_{-jt},W_{-jt}\}}$
Unilateral Bertrand	$\max_{W_{jt}} \Pi_{jt}\big\vert_{\{W_{-jt}\}}$
Unilateral Cournot	$\max_{L_{jt}} \Pi_{jt}\big\vert_{\{L_{-jt}\}}$
Collusive Bertrand of degree $\lambda \in (0, 1]$	$\max_{W_{jt}} \Pi_{jt} + \lambda \sum_{j \neq f} \Pi_{ft} \big\vert_{\{W_{-jt}\}}$
Collusive Cournot of degree $\lambda \in (0, 1]$	$\max_{L_{jt}} \Pi_{jt} + \lambda \sum_{j \neq f} \Pi_{ft} \big\vert_{\{L_{-jt}\}}$

Notes: Each row reports the firm's optimization problem under a candidate conduct. The conditioning sets after the vertical bar indicate which rival decisions the firm takes as given. Under collusive conducts, the firm places weight λ on the sum of rivals' profits. See Appendix Table A1 for the corresponding first-order conditions. For ease of exposition, we suppress the notation that the firm is also choosing M_{jt} to maximize profits, taking K_{jt} as given.

workers are drawn from the outside share, making the size of the outside share especially important under Cournot competition. Under Bertrand oligopsony, the firm chooses its wage taking rivals' wages as given. Rivals' employment adjusts as workers reallocate in response to the wage change of the reference firm, holding wages of competitors fixed. In both cases, a one percent increase in employment at the reference firm leads to greater responses by competitors when the reference firm has a greater market share. An implication is that effective labor supply is less elastic for firms with greater market share, as they anticipate stronger strategic responses from their competitors.

The final two conducts assume that firms collude to jointly suppress wages and employment. Following the IO literature (Bresnahan, 1987; Miller and Weinberg, 2017) as well as the recent labor literature (Delabastita and Rubens, 2025), collusive conduct is modeled as each firm internalizing a weighted sum of its competitors' profits. The parameter $\lambda \in (0, 1]$ dictates the degree of this internalization. In particular, a firm values \$1 of its rivals' total profits equally to $\$\lambda$ of its own profits. In the extreme case in which $\lambda = 1$, the firm is indifferent between \$1 of its own profits and its competitors' profits, which is considered full collusion. Intermediate values, where we utilize $\lambda = 0.5$ in practice, can represent partial collusion.[5] Of course, if firms do not interact strategically (as in monopsonistic competition), collusion is irrelevant, since firms assume their choices do not

[5] An interior value for λ is also a flexible way of capturing other unmodeled reasons for why markdowns exceed Nash-oligopsony levels, such as sequential wage-setting. We refrain from giving a structural interpretation to λ, in line with Corts (1999).

affect competitors' profits, and vice versa. Thus, we can meaningfully define collusion only for the two strategically-interactive conducts: Bertrand and Cournot oligopsonistic competition.

2.4 Wage Markdowns under Each Conduct

The true wage markdown is defined as the ratio of the wage to the marginal revenue product of labor:

$$\Delta_{jt}^{\text{true}} \equiv \frac{W_{jt}}{MRPL_{jt}}.$$

This object is not directly observed because MRPL is unknown. However, for any candidate conduct C from among the eight models described above, the first-order condition of the firm's problem, combined with the residual labor supply function, implies a candidate markdown $\Delta_{jt}^{C}(\boldsymbol{s}_{mt}, \boldsymbol{W}_{mt}, \boldsymbol{\alpha}_m)$ that depends on market shares, wages, and the market-specific labor supply parameters, $\boldsymbol{\alpha}_m$.

The first-order conditions are reported in Appendix Table A1, and those expressions can be rearranged to yield $W_{jt}/MRPL_{jt}$. Under perfect competition, the markdown is simply $\Delta_{jt}^{Perf} = 1$. Under monopsonistic competition, the firm takes all rivals' wages and employment as given. The markdown is

$$\Delta_{jt}^{Mono} = \varepsilon_{jt}^{Mono}/(1+\varepsilon_{jt}^{Mono}), \quad \textit{where} \quad \varepsilon_{jt}^{Mono} \equiv \left.\frac{\partial L_{jt}}{\partial W_{jt}}\right|_{L_{-jt},W_{-jt}} \cdot \frac{W_{jt}}{L_{jt}}.$$

Under unilateral Bertrand oligopsony, the firm takes rivals' wages as given but anticipates that rivals' employment will adjust as workers reallocate. The markdown is

$$\Delta_{jt}^{Bert} = \varepsilon_{jt}^{Bert}/(1+\varepsilon_{jt}^{Bert}), \quad \textit{where} \quad \varepsilon_{jt}^{Bert} \equiv \left.\frac{\partial L_{jt}}{\partial W_{jt}}\right|_{W_{-jt}} \cdot \frac{W_{jt}}{L_{jt}}.$$

Because rivals' employment is free to adjust, $\varepsilon_{jt}^{Bert} < \varepsilon_{jt}^{Mono}$ at the observed wages and employment, implying $\Delta_{jt}^{Bert} < \Delta_{jt}^{Mono}$. Under unilateral Cournot oligopsony, the firm takes rivals' employment as given but anticipates that rivals' wages will adjust to retain their workers, so that effective hiring is only from the outside share. The markdown is

$$\Delta_{jt}^{Cour} = \varepsilon_{jt}^{Cour}/(1+\varepsilon_{jt}^{Cour}), \quad \textit{where} \quad \varepsilon_{jt}^{Cour} \equiv \left(\left.\frac{\partial W_{jt}}{\partial L_{jt}}\right|_{L_{-jt}} \cdot \frac{L_{jt}}{W_{jt}}\right)^{-1}.$$

Because the reference firm assumes it can only attract employment from the outside share, $\varepsilon_{jt}^{Cour} < \varepsilon_{jt}^{Mono}$ at the observed wages and employment, implying $\Delta_{jt}^{Cour} < \Delta_{jt}^{Mono}$.

The collusive conducts are more complex. The first-order conditions involve not only the firm's own supply elasticity but also cross-firm derivatives—$\partial L_{ft}/\partial W_{jt}\big|_{W_{-jt}}$ in the Bertrand case and

$\partial W_{ft}/\partial L_{jt}\big|_{L_{-jt}}$ in the Cournot case—weighted by λ, capturing how the reference firm's choice diverts workers from or bids up wages at competitors. As a result, the collusive markdowns cannot be expressed as a simple function of the firm's own elasticity alone. In Section 4, we consider additively-separable labor preferences with Gumbel-distributed match-specific tastes and show that it is feasible to derive a closed-form expression for the collusive markdown under Cournot competition. However, collusive markdowns must be solved numerically under Bertrand competition.

The key observation is that the candidate conduct C is the true conduct if and only if

$$\Delta_{jt}^{\text{true}} = \Delta_{jt}^{C}(\boldsymbol{s}_{mt}, \boldsymbol{W}_{mt}, \boldsymbol{\alpha}_m), \quad \forall j \in \mathcal{J}_{mt}. \tag{5}$$

The empirical goal is to use the data to reject or accept each candidate conduct, separately for each market, as described in the next section.

3 Conduct Testing

This section develops the identification strategy for conduct testing and describes the Rivers-Vuong pairwise testing procedure.

3.1 Moment Conditions for Conduct Testing

The true markdown $\Delta_{jt}^{\text{true}} \equiv W_{jt}/MRPL_{jt}$ is not observed because MRPL is unknown. Consider a log-additive decomposition of the marginal revenue product:

$$\log MRPL_{jt} = r(\boldsymbol{X}_{jt}, \boldsymbol{\gamma}_m) + \epsilon_{jt}, \tag{6}$$

where the vector of observed determinants of MRPL, $\boldsymbol{X}_{jt}$, is observed; the vector of market-specific parameters that govern the sensitivity of MRPL to observables, $\boldsymbol{\gamma}_m$, is unknown; and ϵ_{jt} is an unobserved residual. In the IO literature on conduct testing, the analogous approach assumes that marginal costs can be decomposed into an unknown function of observables plus an unobserved residual (Backus et al. 2021, Duarte et al. 2024). Substituting the definition of the true markdown, this can be rewritten as

$$\underbrace{\log W_{jt} - r(\boldsymbol{X}_{jt}, \boldsymbol{\gamma}_m)}_{\text{demand-side observable for log markdown}} \quad - \quad \underbrace{\log \Delta_{jt}^{\text{true}}}_{\text{true log markdown}} \quad = \epsilon_{jt}. \tag{7}$$

Suppose we observe an instrumental variable (IV), Z_{jt}^{ϵ}, satisfying three conditions:

Independence. $\mathbb{E}[Z^{\epsilon}_{jt}\epsilon_{jt}|\boldsymbol{X}_{jt}] = 0$.

Exclusion. Z^{ϵ}_{jt} does not belong to $\boldsymbol{X}_{jt}$.

Relevance. $\exists j$ s.t. $\Delta^{\text{true}}_{jt} \neq \Delta^{C}_{jt} \iff \mathbb{E}[Z^{\epsilon}_{jt} \log \Delta^{\text{true}}_{jt}|\boldsymbol{X}_{jt}] \neq \mathbb{E}[Z^{\epsilon}_{jt} \log \Delta^{C}_{jt}|\boldsymbol{X}_{jt}]$.

The first condition states that the IV must be mean independent of the determinants of the log $MRPL$ other than the observables $\boldsymbol{X}_{jt}$. The second condition states that the IV cannot itself be a direct shifter of the log $MRPL$. The third condition is more subtle: it states that, if the conduct is incorrect (in the sense that it implies the wrong markdown for at least one firm), then the covariation between the IV and the (log) true markdown must be different from the covariation between the instrument and the (log) incorrect markdown.

Intuitively, the Relevance condition requires that, in response to variation in the instrument, changes in the supply-side markdown (i.e., the markdown implied by the conduct assumption) should covary differently with changes in the true markdown if the correct conduct is assumed than if an incorrect conduct is assumed, conditional on observables. Unlike in the traditional IV setup, the relevance condition for conduct testing is not directly testable: $\mathbb{E}[Z^{\epsilon}_{jt} \log \Delta^{\text{true}}_{jt}|\boldsymbol{X}_{jt}]$ is unknown, so we cannot check directly if $\mathbb{E}[Z^{\epsilon}_{jt} \log \Delta^{\text{true}}_{jt}|\boldsymbol{X}_{jt}] \neq \mathbb{E}[Z^{\epsilon}_{jt} \log \Delta^{C}_{jt}|\boldsymbol{X}_{jt}]$. However, we can and do verify that $\mathbb{E}[Z^{\epsilon}_{jt} \log \Delta^{C}_{jt}|\boldsymbol{X}_{jt}]$ differs across C, which implies that $\mathbb{E}[Z^{\epsilon}_{jt} \log \Delta^{\text{true}}_{jt}|\boldsymbol{X}_{jt}] = \mathbb{E}[Z^{\epsilon}_{jt} \log \Delta^{C}_{jt}|\boldsymbol{X}_{jt}]$ can hold for at most one candidate conduct C, thus providing a feasible empirical verification that Relevance is satisfied.

The following moment condition is an immediate implication of the three IV conditions:

$$\exists j \text{ s.t. } \Delta^{\text{true}}_{jt} \neq \Delta^{C}_{jt} \iff \mathbb{E}\Big[Z^{\epsilon}_{jt}\Big(\underbrace{\log W_{jt} - r(\boldsymbol{X}_{jt}, \boldsymbol{\gamma}_m)}_{\text{demand-side markdown}}\Big)\Big|\boldsymbol{X}_{jt}\Big] \neq \mathbb{E}\Big[Z^{\epsilon}_{jt}\Big(\underbrace{\log \Delta^{C}_{jt}}_{\text{supply-side markdown}}\Big)\Big|\boldsymbol{X}_{jt}\Big]. \tag{8}$$

This moment condition provides the basis for testing: the proposed conduct C is correct if and only if the supply-side log markdown covaries with the IV to the same extent that the demand-side (proxy for the) log markdown covaries with the IV. This condition has the same intuition as the so-called "Bresnahan (1982) rotation," but due to the availability of a proxy for log $MRPL$, we require only one IV instead of the two IVs required to achieve rotation in the classical approach. See Duarte et al. (2024) and Roussille and Scuderi (2026) for related discussion of the Bresnahan rotation. We discuss our choice of IV in Section 5.

3.2 The Rivers-Vuong Pairwise Test

To compare two candidate conducts C_1 and C_2, we use the Rivers and Vuong (2002) test. Define the GMM objective function, or the "lack of fit" statistic, for a candidate conduct C as

$$Q(C) = g(C)' \mathbb{W}\, g(C),$$

where $g(C) = \mathbb{E}[Z_{jt}(\log W_{jt} - r(X_{jt}, \boldsymbol{\gamma}_m) - \log \Delta_{jt}^{C})]$ is the sample analogue of the identifying moment and $\mathbb{W} = \mathbb{E}[Z_{jt} Z_{jt}']$ is the weighting matrix.

The null hypothesis of the Rivers-Vuong test is that the two conducts fit the data equally well:

$$H_0^{RV} : Q(C_1) = Q(C_2).$$

The test statistic is

$$T_{1,2}^{RV} = \frac{\sqrt{n}(\hat{Q}(C_1) - \hat{Q}(C_2))}{\hat{\sigma}_{RV}}, \tag{9}$$

where $\hat{\sigma}_{RV}$ is an estimate of the standard deviation of the numerator.[6] Under the null, $T_{1,2}^{RV} \sim \mathcal{N}(0,1)$. If $T_{1,2}^{RV} > 1.96$, the test rejects in favor of C_2; if $T_{1,2}^{RV} < -1.96$, the test rejects in favor of C_1. We apply this pairwise test to all pairs of candidate conducts within each market and designate the conduct that achieves the lowest GMM objective as the best-fitting conduct. A conduct is designated as the significant best conduct if it achieves the lowest objective and beats every alternative at the five percent level of significance.

4 A Scalable Empirical Specification

There are three remaining design challenges that we must overcome in order to deploy *at scale* the conduct testing procedure described in the previous section. First, we must specify a practical approach for recovering the relevant components of the production function that allows for diminishing MRPL. It must be feasible in contexts like service industries where it is typically not possible to estimate a full production function, due to the absence of data on capital and materials. Second, we must specify and estimate the labor supply curve such that it is flexible yet still feasible to estimate using a single labor demand shock. Third, given the specification of the labor supply curve, we must be able to solve for analytical (or at least computationally tractable) expressions for the supply-side markdowns across the eight alternative conducts.

[6] The Rivers and Vuong (2002) variance estimator is $\hat{\sigma}_{RV}^2 = 4\left(g(C_1)^2\, \mathbb{W}\, V_{11} + g(C_2)^2\, \mathbb{W}\, V_{22} - 2g(C_1)g(C_2)\, \mathbb{W}\, V_{12}\right)$, where $V_{ck} = \mathbb{E}[\psi_c \psi_k]$ and $\psi_c = \mathbb{W}^{1/2}\left(Z_{jt}(\log W_{jt} - r(X_{jt}, \boldsymbol{\gamma}_m) - \log \Delta_{jt}^{C_c}) - g(C_c)\right) - \frac{1}{2}\mathbb{W}^{3/2} g(C_c)\left(Z_{jt}^2 - 1/\mathbb{W}\right)$.

4.1 Specifying the Marginal Revenue Product

To implement our conduct test, we must specify the predictable component of the log marginal revenue product, $r(\boldsymbol{X}_{jt}, \boldsymbol{\gamma}_{m(j)})$. Consider a Cobb-Douglas revenue production function with heterogeneous total factor productivity (TFP) and market-specific parameters:

$$R_{jt} = A_{jt} L_{jt}^{\gamma_{m(j)}} K_{jt}^{\nu^K_{m(j)}} M_{jt}^{\nu^M_{m(j)}}, \tag{10}$$

where the returns to labor γ_m, returns to capital ν^K_m, and returns to materials ν^M_m may vary across markets. This specification implies $\log MRPL_{jt} = \log \gamma_{m(j)} + \log(R_{jt}/L_{jt}) + \epsilon_{jt}$, where ϵ_{jt} is intended to capture model misspecification error. Conveniently, the $\log MRPL_{jt}$ implied by the Cobb-Douglas production function does not require K_{jt} or M_{jt} to construct.

Following the production function literature, we allow that revenue may be measured with error. Denoting $\hat{R}_{jt}$ as observed revenues, we can write $\hat{R}_{jt} = R_{jt} \exp(-\omega_{jt})$, where ω_{jt} is i.i.d. measurement error. Substituting into the markdown definition yields the key equation we use in the conduct-testing moments:

$$\log W_{jt} - \log(\hat{R}_{jt}/L_{jt}) - \log \Delta^C_{jt} = \log \gamma_{m(j)} + \omega_{jt} + \epsilon_{jt}. \tag{11}$$

Notice that the market-specific returns to labor parameter, γ_m, appears as a market-specific fixed-effect in this equation. Leveraging this equation, we arrive at the modified independence condition corresponding to the conduct test in Section 3: $\mathbb{E}[Z^\epsilon_{jt}(\log \gamma_{m(j)} + \omega_{jt} + \epsilon_{jt})|\boldsymbol{X}_{jt}, m(j)] = 0$.

Our proposed specification $r((\hat{R}_{jt}, L_{jt}), \boldsymbol{\gamma}_{m(j)}) = \log \gamma_{m(j)} + \log(\hat{R}_{jt}/L_{jt})$ has several highly desirable properties while maintaining scalability. First, it accommodates endogenous capital and materials without the need to actually observe either variable, leveraging the Cobb-Douglas functional form. This is important in the context of services industries, for which capital and materials are typically not observed. Second, it permits diminishing MRPL if $\gamma_m < 1$, generalizing from the constant MRPL assumption imposed by Roussille and Scuderi (2026).[7] Third, it does not require that we actually estimate the production function—neither A_{jt}, ν^K_m, nor ν^M_m appears in the relevant equation for testing conduct, while the only production function parameter that does appear, γ_m, can be controlled using additive market-specific fixed-effects. Fourth, it accommodates three sources of unobserved heterogeneity—TFP (A_{jt}), specification error (ϵ_{jt}), and measurement error (ω_{jt})—using the modified independence condition.

[7] Alternatively, decreasing MRPL could be permitted by including additional instruments that shift scale independently of MRPL, as proposed by Duarte et al. (2025) to estimate conduct in product markets with increasing marginal costs. An additional IV is required in their context because they do not directly observe costs—costs in the product market context are analogous to revenues in the labor market context. Since we observe measures of revenues, we are able to test conduct using a single IV.

4.2 Labor Supply Specification and Estimation

The conduct-testing moment condition from Section 3 requires candidate markdowns Δ^C_{jt}, which in turn depend on the residual labor supply curve and its derivatives. We therefore need a labor supply specification that is rich enough to generate distinct markdowns across the eight candidate conducts, parsimonious enough to be estimated using a single source of demand-side variation, and that can be identified without taking a stand on the particular mode of conduct. We structure this estimation in three steps.

Worker Preferences. In practice, we consider a log-separable specification of preferences with an outside option. Worker i's utility from job j is

$$U_{ijt} = \alpha_{m(i)} \log W_{jt} + \xi_{jt} + \eta_{ijt}, \quad \eta_{ijt} \sim \text{Gumbel}(0, 1), \quad \forall j \in \mathcal{J}_{m(i)t}, \tag{12}$$

where we normalize $\alpha_m \log W_{0t|m} + \xi_{0t|m} = 0$ for the outside option, $j = 0$. The parameter α_m governs the sensitivity of workers' choices to wages within market partition m, while ξ_{jt} captures vertical amenities common across workers, and η_{ijt} captures the idiosyncratic match-specific tastes that generate upward-sloping firm-level labor supply. These are the standard preferences from the monopsony literature (Card et al., 2018; Kroft et al., 2025), augmented with a worker-specific outside option, a standard feature in industrial organization.

Two related utility models are worth noting. First, an alternative preference specification common in the labor literature replaces the outside option with a nested logit structure in which workers first choose among local labor markets and then among firms within a market (Lamadon et al., 2022; Berger et al., 2022). In Appendix D, we establish an equivalence result: given any of our conducts, if local nests are small relative to the national economy—as assumed by Berger et al. (2022) and Lamadon et al. (2022)—then markdowns are identical for the nested-logit and logit-with-outside-option specifications. We are not aware of prior work establishing this equivalency. Empirically, we confirm that local nests are sufficiently small as a share of the national workforce that the markdowns differ only up to negligible numerical errors.

Second, in Appendix F, we show that we can test conduct under the more flexible Box-Cox worker utility specification. In particular, in equation (12), the Box-Cox generalization replaces the term $\log W_{jt}$ with the term $\frac{(W_{jt})^{\varrho}-1}{\varrho}$. This term captures the log-wage specification in the special case with $\varrho \to 0$ and the level-wage specification in the special case with $\varrho = 1$, with intermediate values of ϱ spanning the space between the log-linear and linear specifications. Appendix Table A2 demonstrates that this extension allows for substantial heterogeneity in markdowns even under monopsonistic competition. This extension is parsimonious in that it only requires the estimation of one additional parameter, ϱ, using a simple modification of the identifying moment in equation

(13). See Birchall, Mohapatra, and Verboven (2024) for discussion of the Box-Cox specification in product market contexts.

Estimating Wage Sensitivity (α_m). The preference parameter α_m is identified using only within-market variation, and the same estimator recovers α_m regardless of which conduct governs the market. Differencing the labor supply equation across any two firms j and f in the same market eliminates the outside option, which is common to both firms, yielding:

$$(\log s_{ft} - \log s_{jt}) = (\log L_{ft} - \log L_{jt}) = \alpha_m(\log W_{ft} - \log W_{jt}) + (\xi_{ft} - \xi_{jt}),$$

where the first equality follows from the denominator of the market shares being identical for any two firms in the same market. Because wages are correlated with unobserved amenities, consistent estimation requires an instrument Z^{ξ}_{ft} that shifts relative wages through labor demand without directly entering workers' preferences. Given such an instrument satisfying $\mathbb{E}[\xi_{jt} Z^{\xi}_{ft}] = 0, \forall j$, the estimator has a standard instrumental variables representation:

$$\mathbb{E}\left[\left\{(\log L_{ft} - \log L_{jt}) - \alpha_m(\log W_{ft} - \log W_{jt})\right\} Z^{\xi}_{ft}\right] = \mathbb{E}\left[\left\{(\xi_{ft} - \xi_{jt})\right\} Z^{\xi}_{ft}\right] = 0. \tag{13}$$

Importantly, this is one equation in one unknown, allowing for the identification of α_m regardless of the true mode of conduct. The conduct determines *where* on the labor supply curve the market operates, but α_m governs the *shape* of that curve, which is independent of the mode of competition. We describe our choice of the labor demand instrument, Z^{ξ}_{jt}, in Section 5. Conveniently, notice that identification does *not* require that there are no spillover effects of Z^{ξ}_{ft} on L_{jt} or W_{jt}; the only cross-firm restriction is that Z^{ξ}_{ft} is orthogonal to ξ_{jt}.

Recovering the Outside Share ($s_{0t|m}$). The outside share $s_{0t|m}$ is identified from variation across markets, exploiting the fact that α_m has already been recovered from within-market variation. To our knowledge, this is a novel contribution: the existing literature in industrial organization typically calibrates the outside share rather than estimating it.

We partition markets into groups that share common values of α_m and $s_{0t|m}$, and exploit cross-market variation within each group. A positive market-level demand shock raises both wages and total market employment, but the relative magnitude of these responses identifies $s_{0t|m}$. When the outside share is large, there is a deep pool of workers outside the market who can be drawn in, so employment responds strongly relative to wages. When the outside share is small, the market is nearly closed: a demand shock primarily bids up wages among workers who are already employed

inside the market. Using a market-level demand shock ζ_{mt}, we have:

$$s_{0t|m} = \frac{1}{\alpha_m} \frac{\text{Cov}(\log \bar{L}_{mt}, \zeta_{mt})}{\text{Cov}(\bar{w}_{mt}, \zeta_{mt})}, \quad \text{where} \quad \bar{L}_{mt} \equiv \sum_{f \in \mathcal{J}_{mt}} L_{ft}, \quad \bar{w}_{mt} \equiv \sum_{f \in \mathcal{J}_{mt}} s_{ft|m} \log W_{ft}. \tag{14}$$

We derive this expression in Appendix C. As with the estimation of α_m, this estimator is valid regardless of the true conduct because the aggregate relationship between market-level wages and employment is pinned down by the preference structure alone.

4.3 Analytical Expressions for Candidate Markdowns

Given the parametric specification of labor preferences, we solve analytically for the markdowns for each conduct. The results are presented in Table 2 with derivations provided in Appendix B. Key elasticities we derive include $\varepsilon_{jt}^{Mono} = \alpha_m$, $\varepsilon_{jt}^{Bert} = \alpha_m[1 - (1 - s_{0t|m})s_{jt|m}]$, and $\varepsilon_{jt}^{Cour} = \frac{\alpha_m s_{0t|m}}{s_{0t|m} + (1 - s_{0t|m})s_{jt|m}}$. Since $s_{jt} > 0$ implies $(1 - s_{jt}) < 1$ and $(1 + s_{jt}/s_{0t|m}) > 1$, we have $\varepsilon_{jt}^{Cour} < \varepsilon_{jt}^{Bert} < \varepsilon_{jt}^{Mono} = \alpha_m$, with all three converging to α_m as $s_{jt|m} \to 0$. The result that unilateral Cournot markdowns are strictly stronger than unilateral Bertrand markdowns is analogous to the result that Cournot markups are greater than Bertrand markups in the product market (Magnolfi, Quint, Sullivan, and Waldfogel, 2022).

All markdowns can be solved given wages, employment, and $(\alpha_m, s_{0t|m})$ estimates. Computation of the candidate markdowns is straightforward because plug-in expressions are available for perfect competition, monopsonistic competition, unilateral Bertrand, unilateral Cournot, and collusive Cournot. However, collusive Bertrand markdowns, $\Delta_{jt}^{BertColl}$, require a numerical solution for each value of λ_m. Defining $\kappa_m \equiv \sum_{f \in \mathcal{J}_{mt}} (1/\Delta_{ft}^{BertColl} - 1) W_{ft}(1 - s_{0t|m}) s_{ft|m}$, we show that the following system of J_m equations in J_m unknowns provides an implicit solution:

$$\left(\frac{1}{\Delta_{jt}^{BertColl}} - 1 \right) \left[1 - (1 - s_{0t|m}) s_{jt|m} (1 - \lambda_m) \right] = \frac{1}{\alpha_m} + \frac{\lambda_m \kappa_m}{W_{jt}}. \tag{15}$$

Since κ_m depends on the full vector of markdowns, we develop a numerical algorithm that efficiently solves this system of equations.

4.4 Distinguishing Perfect from Monopsonistic Competition in Practice

In general, monopsonistic competition need not imply constant markdowns: under the general preferences of Section 3, the markdown $\varepsilon_{jt}^{Mono}/(1 + \varepsilon_{jt}^{Mono})$ varies across firms with the firm-specific supply elasticity, and so is distinguished from the perfect-competition markdown $\Delta_{jt} = 1$ by its differential correlation with the instrument. However, we adopt the most common empirical

Table 2: Analytical Markdown Expressions by Conduct

Conduct	Markdown Expression
Perfect competition	1
Monopsonistic competition	$\dfrac{\alpha_m}{1+\alpha_m}$
Unilateral Bertrand	$\dfrac{\alpha_m[1-(1-s_{0t\|m})s_{jt\|m}]}{1+\alpha_m[1-(1-s_{0t\|m})s_{jt\|m}]}$
Unilateral Cournot	$\dfrac{\alpha_m}{1+\alpha_m+\frac{1-s_{0t\|m}}{s_{0t\|m}}s_{jt\|m}}$
Collusive Bertrand	requires numerical solution for each λ_m
Collusive Cournot	$\dfrac{\alpha_m}{1+\alpha_m+\frac{1-s_{0t\|m}}{s_{0t\|m}}\left[\lambda_m\frac{\bar{W}_m}{W_{jt}}+(1-\lambda_m)s_{jt\|m}\right]}$ for each λ_m

Notes: This table demonstrates that markdowns can be expressed in terms of the observed (inside) shares, $s_{jt|m} \equiv L_{jt}/\sum_{f\in\mathcal{J}_{mt}} L_{ft}$, observed wages W_{jt}, market average wages $\bar{W}_m \equiv \sum_{f\in\mathcal{J}_{mt}} s_{ft|m}W_{ft}$, and two unobserved market-specific parameters, $(\alpha_m, s_{0t|m})$. We consider $\lambda_m = 1$ (full collusion) and $\lambda_m = 0.5$ (partial collusion). Derivations are provided in Appendix B.

specification of preferences which imply $\varepsilon_{jt}^{Mono} = \alpha_m$ and so the monopsonistic competition markdown is the (market-specific) constant $\alpha_m/(1+\alpha_m)$. Thus, in practice, both perfect and monopsonistic competition imply markdowns that are constant within a market and hence covary identically—that is, not at all—with the IV, implying our main test statistic cannot distinguish them.

Instead, we propose auxiliary tests to discriminate between perfect and monopsonistic competition. First, if we estimate a finite α_m, then residual labor supply slopes upward and every firm has a strict incentive to mark wages down, so perfect competition, which requires $\Delta_{jt} = 1$, cannot be sustained as a Nash equilibrium. Second, if the first stage of the α_m estimation has a large F-statistic, then firm-specific demand shocks move relative wages differentially across firms within a market; under perfect competition all firms are wage-takers and thus do not respond differentially to firm-specific shocks. We confirm both in Section 4.2, leading us to assign all constant markdown findings to the monopsonistic competition category rather than the perfect competition category. This is consistent with the traditional interpretation that an upward-sloping labor supply curve is *de facto* evidence of imperfect competition (Boal and Ransom, 1997; Manning, 2003).

5 Data Sources, Markets, and Instrumental Variables

5.1 Data Sources

Our empirical analysis draws on two establishment-level datasets from the U.S. Census Bureau, each serving a distinct role in the estimation. The first is the Longitudinal Business Database (LBD), which provides annual data on employment and payroll for the near-universe of private-sector employer establishments from 2001 through 2017. LBD establishments are linked to ownership networks using national firm identifiers, which will prove useful below in constructing a shifter of labor demand at the establishment level.

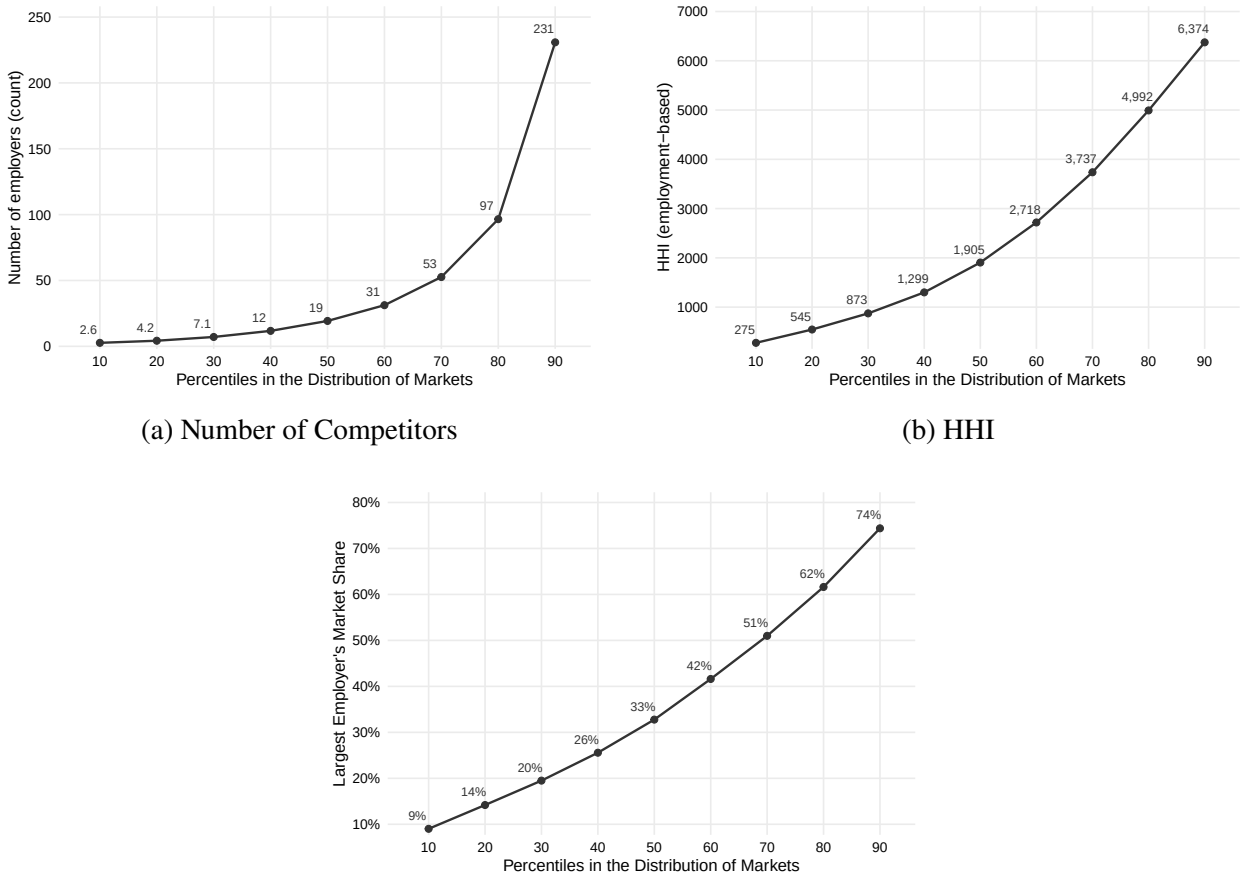
The second source is the quinquennial Economic Census (EC), conducted every five years from 1987 through 2017, which provides establishment-level data on revenues, payroll, and employment. A well-known challenge in U.S. administrative data is that annual revenue reporting is typically at the Employer Identification Number (EIN) level, making it difficult to assign revenues to individual establishments outside of census years. The Economic Census overcomes this limitation by collecting establishment-level revenues directly. However, the EC does not cover all non-farm private-sector industries. We use data from the five sector-specific censuses that report establishment-level revenues and are available to us: Manufacturing, Construction, Retail Trade, Wholesale Trade, and Services. Together, these sectors account for roughly 70 percent of private non-farm employment.[8] The conduct test developed in Section 3 requires establishment-level revenues to construct the demand-side markdown proxy, $\log(\hat{R}_{jt}/L_{jt})$, so we estimate the Rivers-Vuong test using the EC sample, restricting to years in which establishment-level revenues are observed.

5.2 Market Definition

We define a local labor market as a 3-digit NAICS industry by commuting zone pair, similar to Berger et al. (2022), using time-consistent NAICS codes constructed following Fort and Klimek (2018). This yields approximately 108,000 markets. Multi-establishment firms are aggregated to the firm-market-year level.

Figure 2 presents the number of competitors per market, the HHI concentration measure of each market, and the national distribution of the largest inside market share in each market, $s_{jt|m}$, in each market. Subfigure (a) reveals that, in the median market, there are 19 competitors. While this may seem like a large number, subfigure (b) reveals that the median HHI is about 1,900, which is equivalent to about 5 equally-sized competitors. This suggests that markets are dominated by a

[8] We do not observe establishment-level revenues for mining and utilities, FIRE (finance, insurance, and real estate), and the transportation and warehousing sector, but hope to include it in a future update.



(a) Number of Competitors

(b) HHI

(c) Market Share of Largest Employer

Figure 2: Sample Distribution of Market Characteristics

Notes: Markets are equally-weighted.

small number of large competitors. Confirming this intuition, subfigure (c) demonstrates that, in the median market, the largest competitor has 33% market share. Even at the 10th percentile in the distribution of markets, the largest competitor has nearly 10% market share, indicating that nearly all markets have at least one dominant employer. For the purposes of distinguishing oligopsonistic from monopsonistic competition, it is extremely important that there exists at least one firm per market with substantial market share, as the two conducts exhibit nearly identical markdowns among small employers and are thus difficult to distinguish empirically. Thus, it is helpful that even the least-concentrated markets contain an employer with around 10% market share.

The estimation of the labor supply parameters α_m and $s_{0t|m}$ described in Section 4.2 requires pooling across markets. The number of firms per market is often small: highly concentrated industries such as hospitals or dialysis clinics may have only a handful of employers in a given commuting zone, leaving too few observations to estimate α_m separately for each market. Moreover, the between-market estimator of $s_{0t|m}$ exploits cross-market variation in the response of aggregate

employment to aggregate wages, which requires that $s_{0t|m}$ does not vary across the markets being pooled. To gain precision while allowing for meaningful heterogeneity, we partition local markets into 12 broad groups defined by Census region (Northeast, Midwest, South, West) crossed with a three-way industry classification (goods, trade, services). This yields 12 distinct estimates of α_m and $s_{0t|m}$, which are held constant within each partition but allowed to vary across partitions.

5.3 Choices of Instrumental Variables

The estimation requires two instrumental variables (IVs): one for the labor supply estimation (Z^{ξ}_{ft}), which must shift wages through labor demand without affecting workers' preferences, and one for the conduct test (Z^{ϵ}_{ft}), which must satisfy the independence, exclusion, and relevance conditions stated in Section 3. We now describe our chosen IVs.

Labor demand instrument. We propose to use the cross-region housing price shock introduced by Giroud and Mueller (2019) as our labor demand instrument to recover the labor supply parameters, α_m and $s_{0t|m}$. The economic mechanism behind this shock is internal capital market reallocation. When a firm experiences a local housing price shock in one region, the resulting consumer demand (and labor supply) shock in that region affects employment and wages even at locations not directly affected by the shock, given that firms optimally spread the cash-flow impacts across all of its establishments. This cross-region propagation generates establishment-level variation in local labor demand that is driven by conditions in distant markets rather than by local factors. While Giroud and Mueller (2019) study these shocks as objects of interest in their own right, documenting how local demand shocks propagate through firms' internal networks of establishments, we adapt the same variation as an instrumental variable for the labor supply curve.

We measure housing prices using the Zillow Home Value Index at the commuting zone level. The establishment-specific instrument Z^{ξ}_{ft} is the employment-weighted average housing price shock in the other commuting zones in which the firm's same-industry establishments operate. Because different firms within the same local labor market have different multi-establishment networks, they receive different shocks to their labor demand even though they operate in the same market. This within-market variation in establishment-level shocks is what identifies the wage sensitivity parameter α_m: as described in Section 4.2, the estimator differences the labor supply equation across firms j and f in the same market, and the instrument Z^{ξ}_{ft} provides exogenous variation in relative wages across these firms.

To identify the outside share $s_{0t|m}$, we aggregate the establishment-specific shocks to the market

level. The market-wide shock is the share-weighted average of the establishment-specific shocks:

$$\zeta_{mt} \equiv \sum_{j \in \mathcal{J}_{mt}} s_{jt|m} Z^{\xi}_{jt}, \tag{16}$$

where single-establishment firms contribute zero cross-region exposure by construction. As shown in Section 4.2, the outside share is identified from the relative magnitude of the market-level employment and wage responses to this aggregate shock: when the outside share is large, a positive demand shock draws in workers from outside the market, so employment responds strongly relative to wages; when it is small, the same shock primarily bids up wages among incumbent workers. The across-market variation in ζ_{mt}—driven by differences in firms' network structures across markets—provides the exogenous variation needed to estimate this ratio.

The instrument is valid for labor supply estimation if it shifts local wages without directly affecting local workers' preferences over employers. The identifying assumption is that housing price changes in distant regions—which enter the local establishment's labor demand only through the parent firm's internal resource allocation—do not directly affect local workers' tastes for working at that establishment relative to its competitors. This assumption is supported by several features of the design. First, Giroud and Mueller (2019) show that linking establishments to arbitrary distant regions—using equal weights, population weights, income weights, or randomly selected regions—produces elasticities that are indistinguishable from zero, confirming that what matters is the firm's actual network of establishments. Second, the results are robust to saturating the model with highly granular zip code-industry fixed effects, which absorb any local shocks. Third, placebo tests based on counterfactual firm networks—in which establishments are linked to locations where similar firms in the same industry operate, but where the focal firm itself has no presence—yield insignificant estimates, ruling out confounding factors through common clientele or correlated industry demand. These features ensure that the instrument isolates an establishment's demand-side variation originating from the firm's internal network.

Conduct test instrument. For the conduct instrument Z^{ϵ}_{jt}, which enters the Rivers-Vuong testing moment from Section 3, we require a variable satisfying three conditions: independence from the unobserved determinants of the MRPL, exclusion from the MRPL itself, and relevance in the sense that it generates differential covariation with correct and incorrect candidate markdowns.

As our IV, we use the average estimated amenity value among the firm's rivals within its county located in the same market,

$$Z^{\epsilon}_{jt} \equiv \frac{1}{|\mathcal{J}^{c}_{mt} \setminus j|} \sum_{f \in \mathcal{J}^{c}_{mt} \setminus j} \hat{\xi}_{ft}, \tag{17}$$

where $\hat{\xi}_{ft}$ is the amenity recovered from the labor supply estimation and $\mathcal{J}^{c}_{mt} \subset \mathcal{J}_{mt}$ denotes the set

of firms in county c and market m. This instrument has the same logic as the Berry, Levinsohn, and Pakes (1995) instruments widely used in the industrial organization literature (Gandhi and Houde, 2020): characteristics of competing firms shift the competitive environment—and hence the firm's equilibrium markdown—without directly entering the firm's own production function. When rivals become more attractive to workers (higher ξ_{ft}), they draw employment away from firm j, altering the firm's market share and the elasticity of its residual labor supply, shifting markdowns through the competitive channel. The extent to which supply-side markdowns shift in response to rivals' amenities is conduct-dependent, helping to ensure that the relevance condition is satisfied.

A practical advantage of this choice of IV is that it requires no additional data beyond what is already used in the estimation. The amenities $\hat{\xi}_{ft}$ are recovered as residuals from the labor supply estimation, which itself relies only on the five variables assumed throughout: employment, wages, revenues, a market identifier, and the labor demand IV. Once α_m and $s_{0t|m}$ have been estimated using the Giroud-Mueller shock, the rival amenity instrument Z^{ϵ}_{jt} can be constructed directly from the supply-side residuals. This keeps the data requirements of the procedure—supply estimation, markdown construction, and conduct testing—limited to the same five commonly-available variables, maintaining the scalability of our framework.

6 Results

6.1 Labor Supply Parameter Estimates

The first step of our estimation recovers the labor supply parameters using the Giroud-Mueller housing price IV developed in Section 5 and the estimating equation from Section 4.2. Table 3 reports the first-stage F-statistics for each of the 12 market partitions. The F-statistics range from 170 (Trade, Northeast) to 2,503 (Goods, South), all well above conventional thresholds for IV strength. This means that a decline in housing prices in the firm's other locations reduces the wage offered at the focal establishment. These results confirm that a negative financial shock to a firm's collateral value in one region leads the firm to reduce its labor demand and hence its wage offer at establishments in other regions, consistent with Giroud and Mueller (2019). This relationship holds across all 12 partitions of local labor markets, spanning industries and regions.

Given the relevance of the IV, we use the estimator proposed in Section 4 to recover the wage sensitivity parameter α_m, separately within each partition. All 12 estimates are statistically significant at the 5% level, with point estimates ranging from 2.49 (Services, Northeast) to 4.35 (Trade, Northeast). The estimates exhibit a clear pattern across partitions: service industries have the lowest α_m, ranging from 2.49 to 3.09, while goods industries exhibit higher values between 3.27 and 4.07. Trade industries fall in between. Within each industry group, α_m ranges across

regions by 0.80 for goods (3.27 to 4.07), 1.35 for trade (3.00 to 4.35), and 0.60 for services (2.49 to 3.09). Within each region, α_m ranges across industries by 1.86 for the Northeast (2.49 to 4.35), 1.54 for the Midwest (2.53 to 4.07), 0.44 for the South (3.09 to 3.53), and 0.32 for the West (2.95 to 3.27). These magnitudes are consistent with the range of firm-level labor supply elasticity estimates reported in the surveys by Card (2022) and Azar and Marinescu (2024) across empirical contexts and labor demand shifters.

To help interpret these magnitudes, consider the markdowns that we would obtain if every market were governed by monopsonistic competition. Under that conduct, the markdown is the constant $\alpha_m/(1+\alpha_m)$, which depends only on α_m. Evaluated at our estimates, this yields markdowns ranging from 0.71 in the most inelastic partition (Services, Northeast) to 0.81 in the most elastic partition (Trade, Northeast). That is, under monopsonistic competition, workers would receive between 71 and 81 percent of their marginal revenue product, with the gap being largest in service industries where workers' choices are least sensitive to wages. Whether these monopsonistic markdowns accurately describe the data is the question addressed by the conduct test in the sections that follow.

Given the estimates of α_m, we use between-market variation to estimate $s_{0|m}$ for each of the 12 partitions, as described in Section 4.2. The market-level IV regression of log employment on the average log wage, instrumented by the aggregated Giroud-Mueller shock ζ_{mt}, yields a point estimate of 2.025 with a first-stage F-statistic of 7,933, confirming the relevance of the instrument at the market level. Dividing this aggregate elasticity by the within-market α_m estimates yields outside shares ranging from 0.47 (Trade, Northeast) to 0.81 (Services, Northeast), as reported in the final column of Table 3. Service industries, where workers are least sensitive to wages, also have the largest outside shares, while goods and trade industries have substantially smaller outside shares. All estimates lie strictly between zero and one, which the estimator does not impose but which is required by the theory.

6.2 Conduct Distribution

Figure 3 reports the distribution of best-fitting conducts across U.S. labor markets, weighted by employment. Monopsonistic and oligopsonistic competition are by far the most common conducts, accounting for approximately 35 and 40 percent of employment, respectively. Full collusion accounts for 22 percent, while partial collusion accounts for 3 percent. No single conduct dominates, and a substantial share of workers—roughly 25 percent—are in markets best characterized by some form of collusion. In Appendix Figure A1, we further test between Bertrand and Cournot cases among the oligopsonistic and the partially collusive conducts. Within the oligopsonistic category, nearly 60 percent of employment is in markets best fit by Cournot and the remainder by Bertrand, while within the partial collusion category, approximately 43 percent of employment is in markets

Table 3: Labor Supply Parameter Estimates

BLS Super Industry	Census Division	α_m	SE	F-statistic	Implied s_0
Goods	Northeast	3.92	0.13	539	0.52
Goods	Midwest	4.07	0.10	1,131	0.50
Goods	South	3.53	0.06	2,503	0.57
Goods	West	3.27	0.10	951	0.62
Trade	Northeast	4.35	0.42	170	0.47
Trade	Midwest	3.63	0.21	645	0.56
Trade	South	3.50	0.13	1,323	0.58
Trade	West	3.00	0.18	400	0.67
Services	Northeast	2.49	0.09	353	0.81
Services	Midwest	2.53	0.10	802	0.80
Services	South	3.09	0.10	728	0.65
Services	West	2.95	0.14	250	0.69

Notes: Each row reports estimates for one of the 12 market partitions defined by broad industry group (goods, trade, services) crossed with Census region (Northeast, Midwest, South, West). The parameter α_m is the log-wage sensitivity of labor supply, estimated using the IV moment condition in Section 4.2. The instrument is the Giroud-Mueller cross-region housing price shock described in Section 5.3. SE denotes the standard error of $\hat{\alpha}_m$. The F-statistic refers to the strength of the instrument in the first-stage.

best fit by Cournot and the remainder by Bertrand.

The light grey bars in Figure 3 restrict to markets in which the best-fitting conduct is statistically significantly better than all alternatives at the 5 percent level. The proportions decrease slightly, as some null hypotheses cannot be rejected, but the qualitative pattern is preserved.

6.3 Markdown Distribution by Conduct

Figure 4 compares the distribution of wage markdowns under our preferred specification, which allows each market's conduct to be determined by the data, to the distributions implied by assuming each of the eight candidate conducts uniformly across all markets. The preferred specification produces markdowns that are both lower on average and more dispersed than any of the unilateral conduct assumptions. The median markdown under the preferred specification is approximately 72 percent of MRPL, compared to approximately 77 percent under monopsonistic competition, 75 percent under unilateral Bertrand oligopsony, and 75 percent under unilateral Cournot oligopsony. The 10th–90th percentile range under the preferred specification spans nearly 15 percentage points, compared to 8 to 9 percentage points under the unilateral conducts.

The greater dispersion arises because the preferred specification blends markets operating under different conducts, and the collusive conducts generate markdowns that are substantially

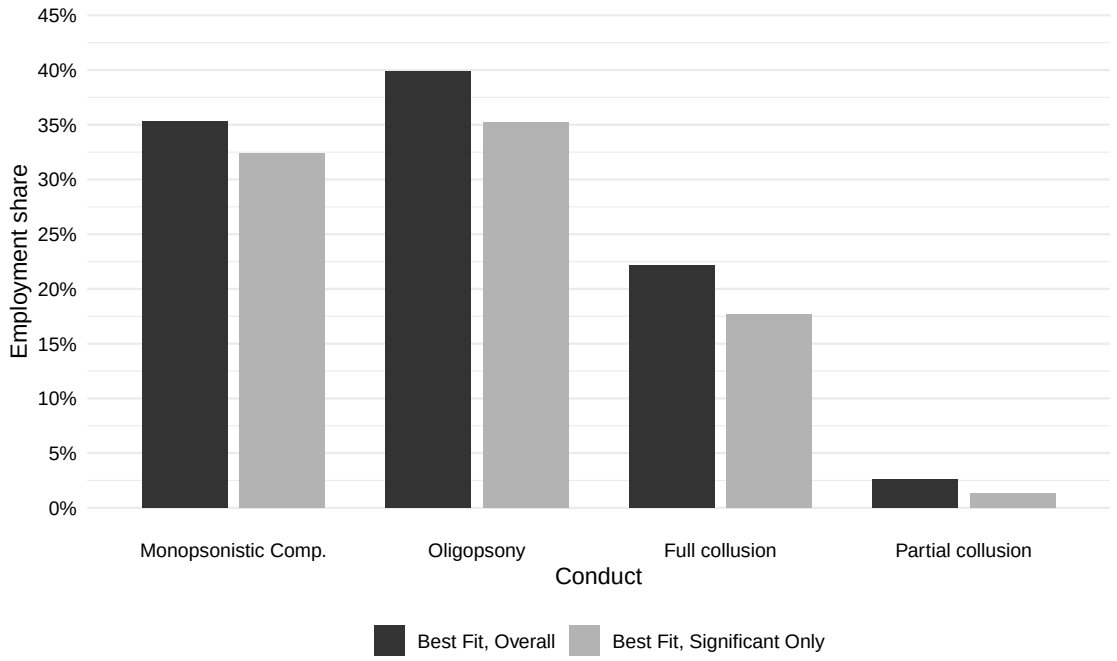


Figure 3: Best Conduct and Significant Best Conduct by Share of Employment

Notes: Each bar reports the share of employment in markets for which the indicated conduct achieves the lowest GMM objective in the Rivers-Vuong testing procedure described in Section 3.2. Dark bars include all markets; light grey bars restrict to markets in which the best-fitting conduct is statistically significant at the 5% level. The oligopsonistic and partial collusion categories pool Bertrand and Cournot subcategories; see Appendix Figure A1 for the Bertrand versus Cournot comparison. Markets are defined as 3-digit NAICS industry by commuting zone pairs.

lower on average and have markedly longer left tails than the unilateral conducts. Imposing full collusion uniformly yields a median markdown of approximately 67 percent, while partial collusion yields markdowns of around 71 percent in the Bertrand case and 70 percent in the Cournot case. Markets classified as collusive under the preferred specification therefore pull the left tail of the overall distribution well below what any unilateral conduct would imply. Meanwhile, the roughly 75 percent of employment in markets classified as monopsonistic or oligopsonistic contributes markdowns in the 70–80 percent range.

Within the non-collusive conducts, nearly all variation in markdowns is driven by heterogeneity in the labor preference parameter α_m across the 12 broad market partitions rather than by cross-firm differences in market shares. This is because only 1 percent of employers have a within-market employment share of 10 percent or more, which is the range where oligopsonistic markdowns diverge meaningfully from the monopsonistic benchmark. For the vast majority of firms, market shares are small enough that the $\alpha_m/(1+\alpha_m)$ approximation from monopsonistic competition applies nearly exactly, and the Bertrand and Cournot oligopsony markdowns are virtually indistinguishable

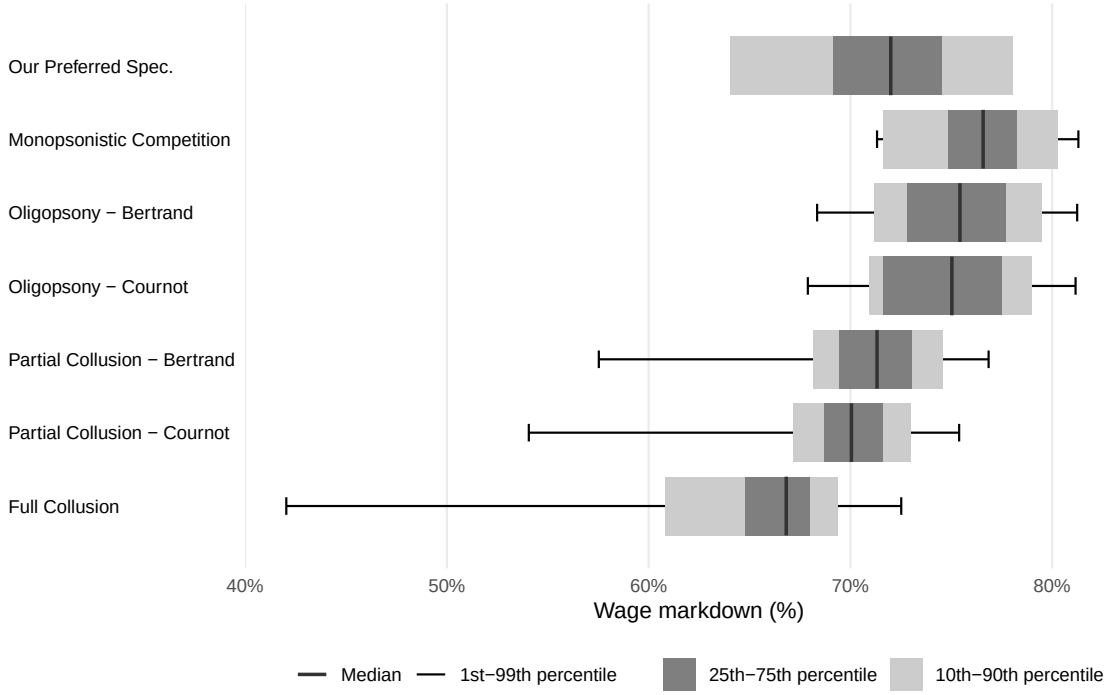


Figure 4: Quantiles in the Markdown Distribution: Estimated vs. Assumed

Notes: Each row displays the distribution of firm-level wage markdowns ($W_{jt}/MRPL_{jt}$, expressed as a percentage) under a given conduct assumption. For each of the eight candidate conducts, markdowns are computed for every firm in every market using the analytical expressions in Table 2 and the estimated labor supply parameters ($\hat{\alpha}_m, \hat{s}_{0|m}$). The first row, labeled "Our Preferred Spec.", assigns each market the conduct that achieves the lowest GMM objective in the Rivers-Vuong testing procedure (Section 3.2), so that different markets may operate under different conducts. All other rows impose the indicated conduct uniformly across all markets. Within each row, the dark vertical line marks the median, the darker shaded region spans the 25th to 75th percentiles, the lighter shaded region spans the 10th to 90th percentiles, and the whiskers extend to the 1st and 99th percentiles. All statistics are weighted by employment.

from each other and from monopsonistic competition. This explains why the three unilateral conduct rows in Figure 4 are very similar, except at the lowest quantiles, which correspond to the highest market shares.

This finding helps resolve a puzzle in the literature. Demand-side markdowns, such as those estimated by Yeh et al. (2022) using the ratio of the wage to a production-function estimate of MRPL, tend to be lower and more dispersed than supply-side markdowns computed under a single national conduct assumption. Our results suggest that the discrepancy arises not from misspecification of the production function or the labor supply curve, but from the imposition of a uniform conduct across heterogeneous markets. Once we allow conduct to vary across markets—and in particular allow for full or partial collusion in the nearly 25 percent of employment where the data support it—the supply-side markdown distribution aligns with the demand-side distribution.

Table 4: Market Characteristics Predicting Collusion

Outcome	Effect on Collusion Probability	
Log Number of Competitors	-0.0153	(0.0018)
Entry rate	-0.0573	(0.0151)
Exit rate	-0.0387	(0.0153)
Variance of MRPL	-0.0398	(0.0024)
Variance of amenities	-0.0006	(0.0003)
N	226,000	

Notes: Each column reports OLS estimates from a linear probability model. The dependent variable equals one if the market's best-fitting conduct is collusive in the Rivers-Vuong test (Section 3.2). Entry rate indicates the ratio of entering establishments to incumbent establishments; exit rate is defined analogously. Variance of MRPL is the within-market variance of $\log(\hat{R}_{jt}/L_{jt})$; variance of amenities is the within-market variance of $\hat{\xi}_{jt}$. An observation is a market-year. Standard errors are in parentheses.

7 Antitrust Implications

7.1 Predictors of Collusion

Classical industrial organization theory offers several predictions about which market characteristics facilitate the sustainability of collusion. Here, we consider whether or not these predictions can be translated into labor market analogues.

The first prediction is that collusion is easier to sustain when there are fewer firms, because coordination is simpler and deviations are more easily detected (Stigler, 1964). The labor-side analogue is the log of the number of establishments in the market. The second prediction is that collusion is more stable when barriers to entry are high, because new entrants would erode the cartel's rents (Bain, 1956; Stigler, 1964). We proxy entry barriers using the ratio of entering establishments to incumbent establishments. The third prediction is that collusion is harder to sustain when firms differ substantially in their costs, because low-cost firms have stronger incentives to deviate from the collusive agreement (Scherer, 1980; Rotemberg and Saloner, 1986). The labor-side analogue is the variance of the marginal revenue product of labor across establishments in the market. The fourth prediction concerns product differentiation, whose effect on collusion is theoretically ambiguous. On one hand, differentiation can soften competition and make collusion less necessary; on the other, it can make deviations harder to detect (Deneckere, 1983; Ross, 1992). The labor-side analogue is the variance of estimated amenities across establishments in the market.

Motivated by the theoretical predictions, Table 4 reports regressions of an indicator for whether a market is classified as significantly collusive on market characteristics. First, we find that markets with more competitors are significantly less likely to be collusive. Second, we find that markets with

Table 5: DOJ Labor Antitrust Cases

Case	Industry	Status	Collusion?
USA v. Surgical Care Affiliates	Surgery	Dismissed	Yes (insign.)
USA v. Patel	Aircraft engines	Acquitted	No
USA v. DaVita	Dialysis clinics	Acquitted	Yes
USA v. Ryan Hee and VDA OC	Nursing	Guilty Plea	Yes
USA v. Jindal	Healthcare staffing	Acquitted	Yes
USA v. Manahe	Home health agencies	Acquitted	No
USA v. Lopez	Nursing	Convicted	Yes

Notes: This table includes all criminal labor market collusion cases filed by the DOJ Antitrust Division, drawing on the list in Posner (2024) and supplemented with subsequent cases from the DOJ website (accessed November 2025). The 'Status' column reports the legal disposition. The 'Collusion?' column reports whether our Rivers-Vuong conduct test (Section 3.2) classifies the corresponding 3-digit NAICS × commuting zone market as collusive: 'Yes' if a collusive conduct achieves the lowest GMM objective and significantly beats all non-collusive alternatives at the 5% level; 'Yes (insign.)' if collusion fits best but does not achieve significance; 'No' if a non-collusive conduct fits best. Collusion is defined in the economic and not legal sense.

higher entry rates and exit rates are significantly less likely to exhibit collusion, consistent with the prediction that barriers to entry protect collusive arrangements. Third, markets with higher variance of the marginal revenue product of labor or amenities are less likely to be collusive, analogous to the prediction that cost heterogeneity makes coordination less likely. Fourth, the variance of amenities has a small negative coefficient, suggesting that greater preference-side differentiation makes coordination less likely. This confirms these key predictions of classical IO theory, but in the labor market: markets are less likely to be collusive when there are more competitors, lower barriers to entry, and more differentiation.

7.2 DOJ Case Studies

As a form of external validation, we apply the conduct test to seven specific labor markets that have been investigated by the Department of Justice for employer collusion. We construct a list of all criminal labor collusion cases filed by the DOJ.[9] Table 5 reports the results. Our test detects collusion in 5 of the 7 cases, a rate of 71.4 percent. Among the two cases that resulted in either a conviction or guilty plea, our test detects collusion in both. By comparison, across all markets in our sample, the baseline rate of detected collusion is 25 percent.

In USA v. Ryan Hee and VDA OC and USA v. Lopez, both involving nursing, the DOJ secured a guilty plea and conviction, respectively, and our test detects collusion in both markets.

[9]We assemble this list by including all cases reported by Posner (2024) and all cases reported on the DOJ website using the search queries 'wage-fixing' and 'no-poach', https://www.justice.gov/atr/antitrust-case-filings.

In USA v. DaVita, involving dialysis clinics, and USA v. Jindal, involving healthcare staffing, the defendants were acquitted at trial, but our test detects collusion in the economic sense in both markets. Similarly, in USA v. Surgical Care Affiliates, involving surgery centers, charges were dismissed, whereas our test picks up coordination, though the result is not statistically significant at the 5 percent level. In contrast, in USA v. Patel, involving aircraft engines, and USA v. Manahe, involving home health agencies, our test does not detect collusion.

We emphasize that the conduct test detects collusion strictly in the economic sense—meaning that the data are best explained by a model in which firms internalize the effect of their decisions on competitors' profits—and not in the legal sense of establishing a conspiracy or agreement.

8 Conclusion

We have developed a scalable framework for estimating labor market conduct using only the variables commonly available in administrative and census datasets: wages, employment, revenues, a market indicator, and a labor demand shifter. The framework accounts for a rich menu of conduct alternatives, allows for diminishing MRPL, and is applicable to service industries.

Applying the Rivers-Vuong testing procedure to approximately 108,000 local labor markets in U.S. Census establishment-level micro-data, we find that no single conduct characterizes the economy. Monopsonistic competition best describes 35 percent of employment, oligopsonistic competition 40 percent, full collusion 22 percent, and partial collusion 3 percent. The heterogeneity in conduct matters: the implied markdown distribution under our preferred specification is both lower and more dispersed than what any single uniformly assumed conduct would generate.

Markets exhibiting collusive conduct are characterized by fewer competitors, lower entry and exit rates, and lower dispersion in marginal revenue products and amenities, thus extending the predictions of classical IO theory to the labor market. Applying the test to markets investigated by the Department of Justice for employer collusion, we detect collusion in five of seven criminal cases, including both cases that resulted in a guilty plea or conviction.

References

Angerhofer, T., A. Collard-Wexler, and M. Weinberg (2025). Monopsony and the Countervailing Power of Unions: Evidence from K-12 Teachers. *Working paper*.

Arnold, D. (2025). Mergers and Acquisitions, Local Labor Market Concentration, and Worker Outcomes. *Working Paper*.

Azar, J. and I. Marinescu (2024). Monopsony Power in the Labor Market. In C. Dustmann and T. Lemieux (Eds.), *Handbook of Labor Economics*, Volume 5, pp. 761–827. Elsevier.

Azar, J. A., S. T. Berry, and I. Marinescu (2022). Estimating Labor Market Power. *NBER Working Paper, No. 30365*.

Backus, M., C. Conlon, and M. Sinkinson (2021). Common Ownership and Competition in the Ready-to-Eat Cereal Industry. *NBER Working Paper*.

Bain, J. S. (1956). *Barriers to New Competition*. Cambridge, MA: Harvard University Press.

Berger, D., T. Hasenzagl, K. Herkenhoff, S. Mongey, and E. A. Posner (2025). Merger Guidelines for the Labor Market. *Journal of Monetary Economics 153*, 103785.

Berger, D., K. Herkenhoff, and S. Mongey (2022). Labor Market Power. *American Economic Review 112*(4), 1147–1193.

Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica 63*(4), 841–890.

Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, 242–262.

Berry, S. T. and P. A. Haile (2014). Identification in Differentiated Products Markets Using Market Level Data. *Econometrica 82*(5), 1749–1797.

Birchall, C., D. Mohapatra, and F. Verboven (2024). Estimating substitution patterns and demand curvature in discrete-choice models of product differentiation. *Review of Economics and Statistics*.

BLS (2025). Union Members—2024. `https://www.bls.gov/news.release/union2.htm`. USDL-25-0105. Accessed: 2025-07-09.

Boal, W. M. and M. Ransom (1997). Monopsony in the labor market. *Journal of Economic Literature*.

Bresnahan, T. F. (1982). The Oligopoly Solution Concept Is Identified. *Economics Letters 10*(1-2), 87–92.

Bresnahan, T. F. (1987). Competition and collusion in the american automobile industry: The 1955 price war. *The Journal of Industrial Economics*, 457–482.

Caldwell, S., A. Dube, and S. Naidu (2026). Monopsony Makes it Big. *Journal of Economic Literature*.

Caldwell, S., I. Haegele, and J. Heining (2026). Bargaining and inequality in the labor market. *The Quarterly Journal of Economics 141*(1), 315–371.

Card, D. (2022). Who Set Your Wage? *American Economic Review 112*(4), 1075–1090.

Card, D., A. R. Cardoso, J. Heining, and P. Kline (2018). Firms and Labor Market Inequality: Evidence and Some Theory. *Journal of Labor Economics 36*(S1), S13–S70.

Chan, M., K. Kroft, E. Mattana, and I. Mourifié (2024). An Empirical Framework for Matching with Imperfect Competition. *NBER Working Paper*.

Corts, K. S. (1999). Conduct parameters and the measurement of market power. *Journal of Econometrics 88*(2), 227–250.

Deb, S., J. Eeckhout, A. Patel, and L. Warren (2024). Walras–bowley lecture: Market power and wage inequality. *Econometrica 92*(3), 603–636.

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

Deneckere, R. (1983). Duopoly Supergames with Product Differentiation. *Economics Letters 11*(1-2), 37–42.

Duarte, M., L. Magnolfi, D. Quint, M. Sølvsten, and C. Sullivan (2025). Conduct and Scale Economies: Evaluating Tariffs in the US Automobile Market. *Working Paper*.

Duarte, M., L. Magnolfi, M. Sølvsten, and C. Sullivan (2024). Testing Firm Conduct. *Quantitative Economics 15*(3), 571–606.

Fort, T. C. and S. D. Klimek (2018). The Effects of Industry Classification Changes on US Employment Composition. *Working Paper*.

Gandhi, A. and J.-F. Houde (2020). Measuring Substitution Patterns in Differentiated-Products Industries. *Unpublished manuscript*.

Gibson, M. D. (2025). Employer Market Power in Silicon Valley. *The Economic Journal*.

Giroud, X. and H. M. Mueller (2019). Firms' Internal Networks and Local Economic Shocks. *American Economic Review 109*(10), 3617–3649.

Gottfries, A. and G. Jarosch (2025). Wage Fixing. *NBER Working Paper*.

Herrera-Caicedo, A., J. Jeffers, and E. Prager (2026). Collusion through Common Leadership. *Journal of Political Economy*.

Jarosch, G., J. S. Nimczik, and I. Sorkin (2024). Granular Search, Market Structure, and Wages. *The Review of Economic Studies 91*(6), 3569–3607.

Kline, P. (2025). Labor Market Monopsony: Fundamentals and Frontiers. In C. Dustmann and T. Lemieux (Eds.), *Handbook of Labor Economics*, Volume 6, pp. 655–728. Elsevier.

Kroft, K., Y. Luo, M. Mogstad, and B. Setzler (2025). Imperfect Competition and Rents in Labor and Product Markets: The Case of the Construction Industry. *American Economic Review 115*(9), 2926–2969.

Lachowska, M., A. Mas, R. Saggio, and S. A. Woodbury (2022). Wage posting or wage bargaining? a test using dual jobholders. *Journal of Labor Economics 40*(S1), S469–S493.

Lamadon, T., M. Mogstad, and B. Setzler (2022). Imperfect Competition, Compensating Differentials, and Rent Sharing in the US Labor Market. *American Economic Review 112*(1), 169–212.

Magnolfi, L., D. Quint, C. Sullivan, and S. Waldfogel (2022). Differentiated-Products Cournot Attributes Higher Markups than Bertrand-Nash. *Economics Letters 219*, 110804.

Manning, A. (2003). *Monopsony in Motion: Imperfect Competition in Labor Markets*. Princeton and Oxford: Princeton University Press.

Manning, A. (2021). Monopsony in Labor Markets: A Review. *ILR Review 74*(1), 3–26.

Masur, J. S. and E. A. Posner (2023). Horizontal Collusion and Parallel Wage Setting in Labor Markets. *University of Chicago Law Review 90*(4), 1227–1308.

Miller, N. H. and M. C. Weinberg (2017). Understanding the Price Effects of the MillerCoors Joint Venture. *Econometrica 85*(6), 1763–1791.

Posner, E. A. (2024). The new labor antitrust. *Antitrust LJ 86*, 503–511.

Prager, E. and M. Schmitt (2021). Employer Consolidation and Wages: Evidence from Hospitals. *American Economic Review 111*(2), 397–427.

Rivers, D. and Q. Vuong (2002). Model Selection Tests for Nonlinear Dynamic Models. *The Econometrics Journal 5*(1), 1–39.

Ross, T. W. (1992). Cartel Stability and Product Differentiation. *International Journal of Industrial Organization 10*(1), 1–13.

Rotemberg, J. J. and G. Saloner (1986). A Supergame-Theoretic Model of Price Wars During Booms. *American Economic Review 76*(3), 390–407.

Roussille, N. and B. Scuderi (2026). Bidding for Talent: A Test of Conduct in a High-Wage Labor Market. *American Economic Review*.

Rubens, M. (2023). Market Structure, Oligopsony Power, and Productivity. *American Economic Review 113*(9), 2382–2410.

Scherer, F. M. (1980). *Industrial Market Structure and Economic Performance* (2nd ed.). Chicago: Rand McNally.

Setzler, B. (2025). Labor and Product Market Power, Endogenous Quality, and the Consolidation of the US Hospital Industry. *NBER Working Paper*.

Sharma, G. (2025). Collusion among Employers in India. *Working Paper*.

Stigler, G. J. (1964). A Theory of Oligopoly. *Journal of Political Economy 72*(1), 44–61.

US DOJ and FTC (2023). Merger Guidelines. *Technical Report*.

Yeh, C., C. Macaluso, and B. Hershbein (2022). Monopsony in the US Labor Market. *American Economic Review 112*(7), 2099–2138.

Appendix

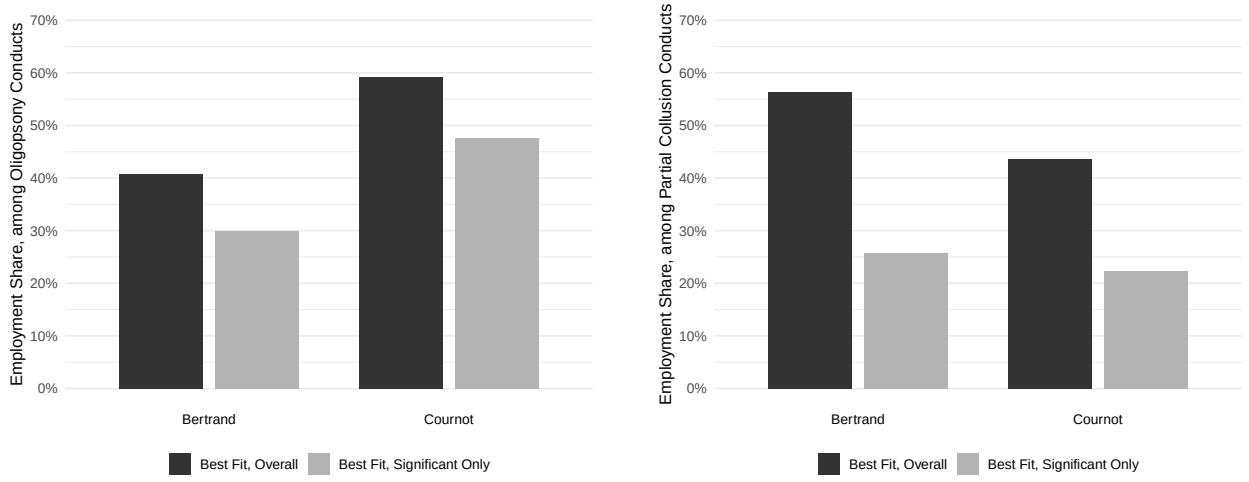
This is the appendix to "Conduct in U.S. Labor Markets," by Michael Rubens, Bradley Setzler, and Chen Yeh, dated August, 2026.

A Additional Tables and Figures

Table A1: First-Order Conditions by Conduct

Conduct	First-Order Condition
Perfect competition	$MRPL_{jt} = W_{jt}$
Monopsonistic competition	$MRPL_{jt} \left.\frac{\partial L_{jt}}{\partial W_{jt}}\right\vert_{L_{-jt}, W_{-jt}} = L_{jt} + W_{jt} \left.\frac{\partial L_{jt}}{\partial W_{jt}}\right\vert_{L_{-jt}, W_{-jt}}$
Unilateral Bertrand	$MRPL_{jt} \left.\frac{\partial L_{jt}}{\partial W_{jt}}\right\vert_{W_{-jt}} = L_{jt} + W_{jt} \left.\frac{\partial L_{jt}}{\partial W_{jt}}\right\vert_{W_{-jt}}$
Unilateral Cournot	$MRPL_{jt} = W_{jt} + L_{jt} \left.\frac{\partial W_{jt}}{\partial L_{jt}}\right\vert_{L_{-jt}}$
Collusive Bertrand of degree λ	$MRPL_{jt} \left.\frac{\partial L_{jt}}{\partial W_{jt}}\right\vert_{W_{-jt}} = L_{jt} + W_{jt} \left.\frac{\partial L_{jt}}{\partial W_{jt}}\right\vert_{W_{-jt}} - \lambda \sum_{f \neq j} (MRPL_{ft} - W_{ft}) \left.\frac{\partial L_{ft}}{\partial W_{jt}}\right\vert_{W_{-jt}}$
Collusive Cournot of degree λ	$MRPL_{jt} = W_{jt} + L_{jt} \left.\frac{\partial W_{jt}}{\partial L_{jt}}\right\vert_{L_{-jt}} + \lambda \sum_{f \neq j} L_{ft} \left.\frac{\partial W_{ft}}{\partial L_{jt}}\right\vert_{L_{-jt}}$

Notes: All derivatives are evaluated at equilibrium values. The Bertrand-type FOCs arise from differentiating profits with respect to the wage, while the Cournot-type FOCs arise from differentiating with respect to employment. The conditioning sets reflect the conduct assumption: under monopsonistic competition, the firm takes both rivals' wages and employment as given; under Bertrand conducts, the firm takes rivals' wages as given; under Cournot conducts, the firm takes rivals' employment as given. In the collusive Bertrand FOC, the cross-firm term involves $\partial L_{ft}/\partial W_{jt}$, which captures the diversion of workers from firm f to firm j following a wage increase at j. The presence of $MRPL_{ft}$ on the right-hand side renders this system implicit, requiring numerical solution. In the collusive Cournot FOC, the cross-firm term involves $\partial W_{ft}/\partial L_{jt}$, which captures the wage increase at rival f when firm j expands employment.



(a) Oligopsony (b) Partial Collusion

Figure A1: Bertrand vs. Cournot Share by Conduct Type

Notes: Each panel reports the share of employment (within the indicated conduct category) that is Bertrand vs. Cournot. Dark bars include all markets; light bars restrict to markets in which the best-fitting conduct is statistically significant at the 5% level.

Table A2: Analytical Markdown Expressions by Conduct under Box-Cox Utility

Conduct	Markdown Expression
Perfect competition	1
Monopsonistic competition	$\frac{W_{jt}^{\varrho}\,\varepsilon_{jt}^{Mono}}{1+W_{jt}^{\varrho}\,\varepsilon_{jt}^{Mono}}$
Unilateral Bertrand	$\frac{W_{jt}^{\varrho}\,\varepsilon_{jt}^{Bert}}{1+W_{jt}^{\varrho}\,\varepsilon_{jt}^{Bert}}$
Unilateral Cournot	$\frac{W_{jt}^{\varrho}\,\varepsilon_{jt}^{Cour}}{1+W_{jt}^{\varrho}\,\varepsilon_{jt}^{Cour}}$
Collusive Bertrand	requires numerical solution for each λ_m
Collusive Cournot	$\frac{\alpha_m W_{jt}^{\varrho}}{1+\alpha_m W_{jt}^{\varrho}+\frac{1-s_{0t\|m}}{s_{0t\|m}}\left[\lambda_m \frac{\bar{W}_m^{(\varrho)}}{W_{jt}^{1-\varrho}}+(1-\lambda_m)s_{jt\|m}\right]}$

Notes: This table generalizes the analytical markdown expressions from Table 2 to the case of Box-Cox worker utility; see Appendix F for derivations. Each unilateral markdown takes the form $W_{jt}^{\varrho}\varepsilon_{jt}^{C}/(1+W_{jt}^{\varrho}\varepsilon_{jt}^{C})$, where the log-case elasticities $\varepsilon_{jt}^{Mono}=\alpha_m$, $\varepsilon_{jt}^{Bert}=\alpha_m[1-(1-s_{0t|m})s_{jt|m}]$, and $(\varepsilon_{jt}^{Cour})^{-1}=\frac{1}{\alpha_m}(1+\frac{1-s_{0t|m}}{s_{0t|m}}s_{jt|m})$ are derived in Appendix B and reported in Section 3; substituting them recovers the fully expanded expressions. Markdowns are expressed in terms of observed inside shares $s_{jt|m}\equiv L_{jt}/\sum_{f\in\mathcal{J}_{mt}}L_{ft}$, observed wages W_{jt}, the Box-Cox market average wage $\bar{W}_m^{(\varrho)}\equiv\sum_{f\in\mathcal{J}_{mt}}s_{ft|m}W_{ft}^{1-\varrho}$, and the market-specific parameters $(\alpha_m, s_{0t|m}, \varrho)$. We consider $\lambda_m=1$ (full collusion) and $\lambda_m=0.5$ (partial collusion). Setting $\varrho=0$ is equivalent to assume that worker utility has the log-wage form considered in the main text, and all markdown expressions become identical to those in Table 2. The collusive Bertrand markdown solves the implicit system in the text for each λ_m.

B Derivation of Supply Elasticities under Logit Preferences

This appendix derives ε_{jt}^{Mono}, ε_{jt}^{Bert}, and ε_{jt}^{Cour} under the preferences specified in Section 4.2, and verifies the markdown expressions in Table 2. Define $\delta_{jt} \equiv \alpha_m \log W_{jt} + \xi_{jt}$ and $\mathcal{W}_{t|m} \equiv \sum_{f \in \mathcal{J}_{mt}} \exp(\delta_{ft})$, so that

$$s_{jt} = \frac{\exp(\delta_{jt})}{1 + \mathcal{W}_{t|m}}, \qquad s_{0t|m} = \frac{1}{1 + \mathcal{W}_{t|m}}.$$

Monopsonistic competition. The firm treats $\mathcal{W}_{t|m}$ as constant, so $\partial s_{jt}/\partial W_{jt} = s_{jt}(\alpha_m/W_{jt})$, giving $\varepsilon_{jt}^{Mono} = \alpha_m$.

Bertrand oligopsony. The firm holds W_{-jt} fixed but allows rival employment to adjust. Applying the quotient rule with $\partial \mathcal{W}_{t|m}/\partial W_{jt} = \exp(\delta_{jt})(\alpha_m/W_{jt})$:

$$\left.\frac{\partial s_{jt}}{\partial W_{jt}}\right|_{W_{-jt}} = s_{jt}(1 - s_{jt})\frac{\alpha_m}{W_{jt}},$$

so $\varepsilon_{jt}^{Bert} = \alpha_m(1 - s_{jt}) = \alpha_m[1 - (1 - s_{0t|m})s_{jt|m}]$.

Cournot oligopsony. The firm holds L_{-jt} (equivalently s_{-jt}) fixed and chooses L_{jt}. From $\log(s_{jt}/s_{0t|m}) = \delta_{jt}$, implicitly differentiate with respect to s_{jt}, using $\partial s_{0t|m}/\partial s_{jt} = -1$:

$$\frac{1}{s_{jt}} + \frac{1}{s_{0t|m}} = \frac{\alpha_m}{W_{jt}}\left.\frac{\partial W_{jt}}{\partial s_{jt}}\right|_{s_{-jt}}.$$

Hence the inverse elasticity is

$$(\varepsilon_{jt}^{Cour})^{-1} = \frac{\partial W_{jt}}{\partial s_{jt}}\frac{s_{jt}}{W_{jt}} = \frac{1}{\alpha_m}\left(1 + \frac{s_{jt}}{s_{0t|m}}\right) = \frac{1}{\alpha_m}\left(1 + \frac{(1 - s_{0t|m})s_{jt|m}}{s_{0t|m}}\right).$$

Collusive Cournot. From the collusive Cournot FOC (Appendix Table A1), the cross-firm term requires $\partial W_{ft}/\partial L_{jt}\big|_{L_{-jt}}$. Holding s_{ft} fixed and differentiating $\log(s_{ft}/s_{0t|m}) = \delta_{ft}$ with respect to s_{jt}:

$$\frac{1}{s_{0t|m}} = \frac{\alpha_m}{W_{ft}}\frac{\partial W_{ft}}{\partial s_{jt}},$$

so $L_{ft}\,\partial W_{ft}/\partial L_{jt} = W_{ft}s_{ft}/(\alpha_m s_{0t|m})$. Substituting both the own and cross terms into the FOC and writing $\sum_{f \neq j} W_{ft}s_{ft} = (1 - s_{0t|m})\bar{W}_m - W_{jt}s_{jt}$ yields the collusive Cournot markdown in Table 2.

Collusive Bertrand. The cross-wage share derivative is $\partial s_{ft}/\partial W_{jt}\big|_{W_{-jt}} = -s_{ft}s_{jt}(\alpha_m/W_{jt})$. Substituting own and cross derivatives into the collusive Bertrand FOC, dividing by $\alpha_m s_{jt}$, and

defining $\kappa_m \equiv \sum_{f \in \mathcal{J}_{mt}} (1/\Delta_{ft} - 1) W_{ft} s_{ft}$:

$$\left(\frac{1}{\Delta_{jt}} - 1\right)\left[1 - s_{jt}(1-\lambda)\right] = \frac{1}{\alpha_m} + \frac{\lambda \kappa_m}{W_{jt}}.$$

Since κ_m depends on the full vector of markdowns, this system requires a numerical solution.

C Identification of the Outside Share

This appendix derives the estimator for $s_{0t|m}$ reported in Section 4.2. The outside share is

$$s_{0t|m} = \frac{1}{1 + \mathcal{W}_{t|m}}, \qquad 1 - s_{0t|m} = \frac{\mathcal{W}_{t|m}}{1 + \mathcal{W}_{t|m}}.$$

Total inside employment is $\bar{L}_{mt} = (1 - s_{0t|m}) \bar{L}^{\text{pot}}_{mt}$, where $\bar{L}^{\text{pot}}_{mt}$ is the (unobserved) number of potential workers. Consider a uniform increase of Δ in all log wages in the market. The wage index becomes

$$\mathcal{W}_{t|m}(\Delta) = \sum_{f \in \mathcal{J}_{mt}} \exp\big(\alpha_m (\log W_{ft} + \Delta) + \xi_{ft}\big) = e^{\alpha_m \Delta} \sum_{f \in \mathcal{J}_{mt}} \exp\big(\alpha_m \log W_{ft} + \xi_{ft}\big) = e^{\alpha_m \Delta} \mathcal{W}_{t|m},$$

so

$$s_{0t|m}(\Delta) = \frac{1}{1 + e^{\alpha_m \Delta} \mathcal{W}_{t|m}}.$$

Differentiating,

$$\left.\frac{\partial s_{0t|m}(\Delta)}{\partial \Delta}\right|_{\Delta=0} = \frac{-\alpha_m \mathcal{W}_{t|m}}{(1 + \mathcal{W}_{t|m})^2} = -\alpha_m \, s_{0t|m}(1 - s_{0t|m}).$$

Since $\log \bar{L}_{mt} = \log(1 - s_{0t|m}) + \log \bar{L}^{\text{pot}}_{mt}$,

$$\left.\frac{\partial \log \bar{L}_{mt}}{\partial \Delta}\right|_{\Delta=0} = \frac{\alpha_m \, s_{0t|m}(1 - s_{0t|m})}{1 - s_{0t|m}} = \alpha_m \, s_{0t|m}.$$

This elasticity is the population coefficient of $\bar{w}_{mt}$ in a regression of $\log \bar{L}_{mt}$ on $\bar{w}_{mt}$, where $\bar{w}_{mt} \equiv \sum_{f \in \mathcal{J}_{mt}} s_{ft|m} \log W_{ft}$. Because $\bar{w}_{mt}$ is endogenous, we instrument with the market-level demand shock ζ_{mt}:

$$\alpha_m \, s_{0t|m} = \frac{\text{Cov}(\log \bar{L}_{mt}, \, \zeta_{mt})}{\text{Cov}(\bar{w}_{mt}, \, \zeta_{mt})}.$$

Dividing by α_m (estimated from within-market variation as in Section 4.2) gives the result.

D Equivalence of Outside-Option Logit and Nested Logit

This appendix shows that, in the small-nest limit, the nested-logit and outside-option specifications yield identical markdowns under every conduct when estimated using the procedure in Section 4.2.

Nested-logit preferences. Replace the passive outside option with cross-market mobility governed by a nesting parameter $\sigma \in [0, 1)$:

$$U_{ijt} = \alpha_m^s \log W_{jt} + \xi_{jt} + \zeta_{imt} + (1-\sigma)\eta_{ijt}, \quad \zeta_{imt} \sim \text{Cardell}(\sigma), \quad \eta_{ijt} \sim \text{Gumbel}(0,1),$$

where α_m^s denotes the structural preference parameter. The within-nest conditional choice probability is

$$s_{jt|m} = \frac{\exp\big((\alpha_m^s \log W_{jt} + \xi_{jt})/(1-\sigma)\big)}{\sum_{f \in \mathcal{J}_{mt}} \exp\big((\alpha_m^s \log W_{ft} + \xi_{ft})/(1-\sigma)\big)},$$

and the nest's share of national employment is $S_{mt} = \mathcal{W}_{t|m}^{1-\sigma} / \sum_{m'} \mathcal{W}_{t|m'}^{1-\sigma}$, where

$$\mathcal{W}_{t|m} \equiv \sum_{f \in \mathcal{J}_{mt}} \exp\big((\alpha_m^s \log W_{ft} + \xi_{ft})/(1-\sigma)\big).$$

What the within-market regression identifies. Differencing across firms within a nest,

$$\log L_{ft} - \log L_{jt} = \frac{\alpha_m^s}{1-\sigma}(\log W_{ft} - \log W_{jt}) + \frac{1}{1-\sigma}(\xi_{ft} - \xi_{jt}),$$

which is similar to the key identifying equation of Berry (1994). The IV estimator in Section 4.2 therefore identifies the composite $\alpha_m \equiv \alpha_m^s/(1-\sigma)$, regardless of the true conduct.

What the aggregate elasticity identifies. Under a uniform increase Δ in all log wages in the nest, the wage index scales as $\mathcal{W}_{t|m}(\Delta) = e^{\alpha_m \Delta} \mathcal{W}_{t|m}$ (by the same argument as Appendix C). The nest share is $S_{mt} = \mathcal{W}_{t|m}^{1-\sigma} / \overline{\mathcal{W}}$, where $\overline{\mathcal{W}} \equiv \sum_{m'} \mathcal{W}_{t|m'}^{1-\sigma}$. Under the uniform wage increase, $\mathcal{W}_{t|m}(\Delta)^{1-\sigma} = e^{\alpha_m(1-\sigma)\Delta} \mathcal{W}_{t|m}^{1-\sigma}$, while all other nests are unaffected. Writing $\overline{\mathcal{W}}_{-m} \equiv \overline{\mathcal{W}} - \mathcal{W}_{t|m}^{1-\sigma}$,

$$S_{mt}(\Delta) = \frac{e^{\alpha_m(1-\sigma)\Delta} \mathcal{W}_{t|m}^{1-\sigma}}{\overline{\mathcal{W}}_{-m} + e^{\alpha_m(1-\sigma)\Delta} \mathcal{W}_{t|m}^{1-\sigma}}.$$

Differentiating and evaluating at $\Delta = 0$:

$$\left.\frac{\partial S_{mt}}{\partial \Delta}\right|_{\Delta=0} = \alpha_m(1-\sigma)\frac{\mathcal{W}_{t|m}^{1-\sigma}\,\overline{\mathcal{W}}_{-m}}{\overline{\mathcal{W}}^2} = \alpha_m(1-\sigma)\, S_{mt}(1-S_{mt}),$$

so $\partial \log S_{mt}/\partial\Delta|_{\Delta=0} = \alpha_m(1-\sigma)(1-S_{mt})$. In the small-nest limit ($S_{mt} \approx 0$), this is $\alpha_m(1-\sigma)$. Applying the outside-share estimator from Section 4.2 returns $\hat{s}_{0t|m} = 1-\sigma$.

Bertrand elasticity. The own-wage labor supply derivative in the nested logit is

$$\frac{\partial s_{jt}}{\partial W_{jt}} = \frac{\alpha_m^s}{(1-\sigma)W_{jt}} s_{jt}[1 - \sigma s_{jt|m} - (1-\sigma)s_{jt}],$$

where $s_{jt} = s_{jt|m} S_{mt}$. In the small-nest limit,

$$\varepsilon_{jt}^{Bert} = \alpha_m(1 - \sigma s_{jt|m}) = \alpha_m\left[1 - (1 - \hat{s}_{0t|m})s_{jt|m}\right],$$

which is the outside-option expression from Appendix B with $s_{0t|m} = 1-\sigma$.

Cournot inverse elasticity. From the nested-logit inverse supply in the small-nest limit,

$$(\varepsilon_{jt}^{Cour})^{-1} = \frac{1}{\alpha_m}\left(1 + \frac{\sigma}{1-\sigma} s_{jt|m}\right) = \frac{1}{\alpha_m}\left(1 + \frac{1-\hat{s}_{0t|m}}{\hat{s}_{0t|m}} s_{jt|m}\right),$$

which is the outside-option expression with $s_{0t|m} = 1-\sigma$.

Monopsonistic competition. Setting $s_{jt|m} = 0$ in either the Bertrand or Cournot expression, both models yield $\varepsilon_{jt}^{Mono} = \alpha_m$.

Collusion: Bertrand cross-effect. The collusive Bertrand FOC (Appendix Table A1), divided by L_{jt}, depends on the own and cross ratios $\frac{1}{s_{jt}}\frac{\partial s_{jt}}{\partial W_{jt}}$ and $\frac{1}{s_{jt}}\frac{\partial s_{ft}}{\partial W_{jt}}$. The own ratio was shown above. For the cross ratio, in the outside-option model:

$$\frac{1}{s_{jt}}\frac{\partial s_{ft}}{\partial W_{jt}}\bigg|_{W_{-jt}} = -s_{ft}\frac{\alpha_m}{W_{jt}} = -(1 - s_{0t|m})s_{ft|m}\frac{\alpha_m}{W_{jt}}.$$

In the nested logit, the small-nest cross-derivative (from $\partial s_{ft}/\partial W_{jt} = -\frac{\alpha_m^s}{(1-\sigma)W_{jt}} s_{jt}[\sigma s_{ft|m} + (1-\sigma)s_{ft}]$ with $s_{ft} \approx 0$) gives

$$\frac{1}{s_{jt}}\frac{\partial s_{ft}}{\partial W_{jt}}\bigg|_{W_{-jt}} = -\frac{\alpha_m^s \sigma s_{ft|m}}{(1-\sigma)W_{jt}} = -\sigma s_{ft|m}\frac{\alpha_m}{W_{jt}}.$$

Under $s_{0t|m} = 1-\sigma$, these are identical. Since both the own and cross ratios match, the collusive Bertrand FOC—and hence its numerical solution for Δ_{jt}—is the same under both specifications.

Collusion: Cournot cross-effect. The collusive Cournot FOC (Appendix Table A1) involves $L_{ft}\,\partial W_{ft}/\partial L_{jt}\big|_{L_{-jt}}$. In the outside-option model, from Appendix B:

$$L_{ft}\frac{\partial W_{ft}}{\partial L_{jt}}\bigg|_{L_{-jt}} = \frac{W_{ft}\,s_{ft}}{\alpha_m\,s_{0t|m}} = \frac{W_{ft}\,(1-s_{0t|m})\,s_{ft|m}}{\alpha_m\,s_{0t|m}}.$$

In the nested logit, the inverse supply for firm f in the small-nest limit is (from Appendix B, applied to the nested-logit shares):

$$\log W_{ft} = \frac{1}{\alpha_m}\log L_{ft} + \frac{\sigma}{\alpha_m(1-\sigma)}\log \bar{L}_{mt} + \text{terms not involving } L_{jt},$$

where $\bar{L}_{mt} = \sum_{k\in\mathcal{J}_{mt}} L_{kt}$. Holding L_{ft} and all L_{kt} for $k \neq j$ fixed, $\partial \log \bar{L}_{mt}/\partial \log L_{jt} = L_{jt}/\bar{L}_{mt} = s_{jt|m}$, so

$$\frac{\partial \log W_{ft}}{\partial \log L_{jt}}\bigg|_{L_{-jt}} = \frac{\sigma}{\alpha_m(1-\sigma)}s_{jt|m}.$$

Converting to levels, $L_{ft}\,\partial W_{ft}/\partial L_{jt} = W_{ft}(L_{ft}/L_{jt})\cdot\frac{\sigma\,s_{jt|m}}{\alpha_m(1-\sigma)} = \frac{W_{ft}\,s_{ft|m}\,\sigma}{\alpha_m(1-\sigma)}$, where the last step uses $L_{ft}/L_{jt} = s_{ft|m}/s_{jt|m}$. Under $s_{0t|m} = 1-\sigma$, this equals the outside-option expression. Since both the own inverse elasticity and the cross-effect match, the collusive Cournot markdown is identical under both specifications.

E Conduct Testing with Heterogeneous Worker Types

The main text treats labor within a market as homogeneous. This appendix extends the framework to allow for discrete worker types, indexed by $\chi \in \mathcal{X}$. Examples of data that may be used to define worker types include educational attainment, occupation, and worker fixed effects. Each type has its own labor supply curve, its own markdown, and—within a given market—its own conduct. The test statistic from Section 3 applies separately by type.

Worker preferences and labor supply by type. Let L^{χ}_{jt} denote establishment j's employment of type-χ workers and W^{χ}_{jt} the corresponding wage. Workers of type χ have preferences over jobs as in Section 4.2, with type-specific parameters:

$$U^{\chi}_{ijt} = \alpha^{\chi}_{m(i)}\log W^{\chi}_{jt} + \xi^{\chi}_{jt} + \eta^{\chi}_{ijt}, \quad \eta^{\chi}_{ijt} \sim \text{Gumbel}(0,1), \quad \forall j \in \mathcal{J}_{m(i)t}.$$

Each type chooses among jobs independently, so the type-χ inside share is $s^{\chi}_{jt|m} \equiv L^{\chi}_{jt}/\sum_{f\in\mathcal{J}_{mt}} L^{\chi}_{ft}$, with outside share $s^{\chi}_{0t|m}$ defined analogously.

Estimating α_m^χ. The wage-sensitivity parameter is recovered within type by the procedure of Section 4.2. Differencing the type-χ labor supply equation across establishments j and f in the same market eliminates the outside option, and instrumenting relative wages with the labor demand shock Z_{ft}^ξ of Section 5.3 gives

$$\mathbb{E}\left[\left\{(\log L_{ft}^\chi - \log L_{jt}^\chi) - \alpha_m^\chi(\log W_{ft}^\chi - \log W_{jt}^\chi)\right\} Z_{ft}^\xi\right] = 0, \quad \forall \chi \in \mathcal{X}.$$

This is one equation in the one unknown α_m^χ, identified regardless of the true conduct. The outside share $s_{0t|m}^\chi$ is then recovered from between-market variation as in Section 4.2. The residuals of the type-χ supply estimation yield type-specific amenities $\hat{\xi}_{ft}^\chi$.

Revenue production. Establishment revenue is Cobb-Douglas in the type-specific labor inputs, capital, and materials:

$$R_{jt} = A_{jt}\left(\prod_{\chi\in\mathcal{X}} \left(L_{jt}^\chi\right)^{\gamma_{m(j)}^\chi}\right) K_{jt}^{\nu_{m(j)}^K} M_{jt}^{\nu_{m(j)}^M}.$$

The marginal revenue product of type-χ labor is

$$MRPL_{jt}^\chi \equiv \frac{\partial F_{jt}}{\partial L_{jt}^\chi} = \gamma_{m(j)}^\chi \frac{R_{jt}}{L_{jt}^\chi},$$

so that $\log MRPL_{jt}^\chi = \log \gamma_{m(j)}^\chi + \log(R_{jt}/L_{jt}^\chi) + \epsilon_{jt}^\chi$. This is the same expression as in Section 4.1, with the type-χ labor input replacing total labor and γ_m^χ replacing γ_m. As before, neither A_{jt}, K_{jt}, nor M_{jt} appears, and γ_m^χ enters as a market-by-type fixed effect.

Markdowns and the test. The type-χ markdown is $\Delta_{jt}^{\text{true},\chi} \equiv W_{jt}^\chi / MRPL_{jt}^\chi$. For each candidate conduct C, the first-order conditions in Appendix Table A1—taken with respect to W_{jt}^χ or L_{jt}^χ, holding the other types fixed—yield a candidate markdown $\Delta_{jt}^{C,\chi}$. Because the production function is separable in the type inputs and each type's supply depends only on its own wages, the analytical expressions in Table 2 carry over verbatim with $(\alpha_m^\chi, s_{0t|m}^\chi, s_{jt|m}^\chi)$ in place of $(\alpha_m, s_{0t|m}, s_{jt|m})$.

Substituting the revenue specification into the markdown definition gives the type-χ analogue of the key estimating equation:

$$\log W_{jt}^\chi - \log(\hat{R}_{jt}/L_{jt}^\chi) - \log \Delta_{jt}^{C,\chi} = \log \gamma_{m(j)}^\chi + \omega_{jt} + \epsilon_{jt}^\chi.$$

The conduct instrument need not be type-specific: the same rival-amenity instrument Z_{jt}^ϵ of

Section 5.3 is reused across types. The type-χ lack-of-fit statistic is

$$Q^{\chi}(C) = g^{\chi}(C)'\mathbb{W}\, g^{\chi}(C), \quad g^{\chi}(C) = \mathbb{E}\big[Z^{\epsilon}_{jt}\big(\log W^{\chi}_{jt} - \log(\hat{R}_{jt}/L^{\chi}_{jt}) - \log \Delta^{C,\chi}_{jt}\big)\big],$$

and the Rivers-Vuong statistic comparing conducts C_1 and C_2 for type χ is

$$T^{RV,\chi}_{1,2} = \frac{\sqrt{n}\big(\hat{Q}^{\chi}(C_1) - \hat{Q}^{\chi}(C_2)\big)}{\hat{\sigma}^{\chi}_{RV}}, \quad \forall \chi \in \mathcal{X}.$$

The best-fitting conduct is selected market-by-type. Two types in the same market may therefore operate under different conducts. For example, firms may collude over the wages of non-college-educated workers yet compete unilaterally over college-educated workers.

F Conduct Testing under Box-Cox Worker Utility

This appendix derives conduct-specific markdowns and identification under the Box-Cox generalization of the worker utility in equation (12). In this extension, worker i's utility from job j is

$$U_{ijt} = \delta_{jt} + \eta_{ijt}, \quad \eta_{ijt} \sim \text{Gumbel}(0,1), \quad \forall j \in \mathcal{J}_{m(i)t},$$

where,

$$\delta_{jt} = \alpha_m \frac{W^{\varrho}_{jt} - 1}{\varrho} + \xi_{jt}, \qquad \delta_{0t|m} = 0,$$

with $s_{jt} = \exp(\delta_{jt})/(1+\mathcal{W}_{t|m})$ and $s_{0t|m} = 1/(1+\mathcal{W}_{t|m})$ as before, where $\mathcal{W}_{t|m} \equiv \sum_{f\in\mathcal{J}_{mt}} \exp(\delta_{ft})$. The logit share structure is unchanged; the key difference is the marginal utility of the wage,

$$\frac{\partial \delta_{jt}}{\partial W_{jt}} = \alpha_m W^{\varrho-1}_{jt}.$$

We now establish that the Box-Cox transform rescales each unilateral supply elasticity by the firm's own wage raised to the ϱ power,

$$\varepsilon^{(\varrho),C}_{jt} = W^{\varrho}_{jt}\, \varepsilon^{C}_{jt}, \qquad C \in \{Mono, Bert, Cour\},$$

where ε^{C}_{jt} is the corresponding log-case elasticity of Appendix B. The markdowns then follow by substitution, and every expression collapses to its log-wage counterpart as $\varrho \to 0$.

Monopsonistic competition. The firm treats $\mathcal{W}_{t|m}$ as constant, so $\partial s_{jt}/\partial W_{jt} = s_{jt}\,\alpha_m W^{\varrho-1}_{jt}$, giving $\varepsilon^{(\varrho),Mono}_{jt} = \alpha_m W^{\varrho}_{jt} = W^{\varrho}_{jt}\,\varepsilon^{Mono}_{jt}$. Unlike the log case, the elasticity is no longer the market

constant α_m: it scales with the wage level, so the markdown varies across firms within a market.

Bertrand oligopsony. Holding W_{-jt} fixed and allowing rival employment to adjust, the quotient rule with $\partial \mathcal{W}_{t|m}/\partial W_{jt} = \exp(\delta_{jt})\,\alpha_m W_{jt}^{\varrho-1}$ gives

$$\left.\frac{\partial s_{jt}}{\partial W_{jt}}\right|_{W_{-jt}} = s_{jt}(1-s_{jt})\,\alpha_m W_{jt}^{\varrho-1},$$

so $\varepsilon_{jt}^{(\varrho),Bert} = \alpha_m W_{jt}^{\varrho}(1-s_{jt}) = \alpha_m W_{jt}^{\varrho}[1-(1-s_{0t|m})s_{jt|m}] = W_{jt}^{\varrho}\,\varepsilon_{jt}^{Bert}$.

Cournot oligopsony. Holding L_{-jt} (equivalently s_{-jt}) fixed, invert the share equation. From $\log(s_{jt}/s_{0t|m}) = \delta_{jt}$, implicit differentiation with respect to s_{jt}, using $\partial s_{0t|m}/\partial s_{jt} = -1$ and $\partial \delta_{jt}/\partial W_{jt} = \alpha_m W_{jt}^{\varrho-1}$, gives

$$\frac{1}{s_{jt}} + \frac{1}{s_{0t|m}} = \alpha_m W_{jt}^{\varrho-1}\left.\frac{\partial W_{jt}}{\partial s_{jt}}\right|_{s_{-jt}},$$

so the inverse elasticity is

$$(\varepsilon_{jt}^{(\varrho),Cour})^{-1} = \frac{\partial W_{jt}}{\partial s_{jt}}\frac{s_{jt}}{W_{jt}} = \frac{1}{\alpha_m W_{jt}^{\varrho}}\left(1+\frac{s_{jt}}{s_{0t|m}}\right) = \frac{1}{\alpha_m W_{jt}^{\varrho}}\left(1+\frac{(1-s_{0t|m})s_{jt|m}}{s_{0t|m}}\right),$$

so that $\varepsilon_{jt}^{(\varrho),Cour} = W_{jt}^{\varrho}\,\varepsilon_{jt}^{Cour}$.

Collusive Cournot. The collusive Cournot FOC (Appendix Table A1) requires the cross term $\partial W_{ft}/\partial L_{jt}\big|_{L_{-jt}}$. Holding s_{ft} fixed and differentiating $\log(s_{ft}/s_{0t|m}) = \delta_{ft}$ with respect to s_{jt},

$$\frac{1}{s_{0t|m}} = \alpha_m W_{ft}^{\varrho-1}\frac{\partial W_{ft}}{\partial s_{jt}},$$

so $L_{ft}\,\partial W_{ft}/\partial L_{jt} = W_{ft}^{1-\varrho}\,s_{ft}/(\alpha_m s_{0t|m})$. Substituting the own and cross terms into the FOC and defining the Box-Cox market average wage

$$\bar{W}_m^{(\varrho)} \equiv \sum_{f\in\mathcal{J}_{mt}} s_{ft|m}\,W_{ft}^{1-\varrho},$$

so that $\sum_{f\neq j} s_{ft} W_{ft}^{1-\varrho} = (1 - s_{0t|m}) W_m^{(\varrho)} - s_{jt} W_{jt}^{1-\varrho}$, and collecting terms gives

$$\Delta_{jt}^{(\varrho),CourColl} = \frac{\alpha_m W_{jt}^{\varrho}}{1 + \alpha_m W_{jt}^{\varrho} + \frac{1-s_{0t|m}}{s_{0t|m}} \left[(1-\lambda_m) s_{jt|m} + \lambda_m \frac{\bar{W}_m^{(\varrho)}}{W_{jt}^{1-\varrho}} \right]}.$$

The wage enters the collusive term through $\bar{W}_m^{(\varrho)}/W_{jt}^{1-\varrho}$ rather than $\bar{W}_m/W_{jt}$; unlike the elasticities, the market average wage index $\bar{W}_m^{(\varrho)}$ is not a scalar multiple of its log-case counterpart $\bar{W}_m$, since the exponent sits inside the sum. Setting $\lambda_m = 0$ recovers $\Delta_{jt}^{(\varrho),Cour}$; setting $\varrho = 0$ gives $W_{jt}^{\varrho} = 1$ and $W_m^{(\varrho)} = \bar{W}_m$, recovering the log-utility expression in Table 2. All Cournot markdowns remain available in closed form.

Collusive Bertrand. The cross-wage share derivative is $\partial s_{ft}/\partial W_{jt}\big|_{W_{-jt}} = -s_{ft} s_{jt}\, \alpha_m W_{jt}^{\varrho-1}$. Substituting own and cross derivatives into the collusive Bertrand FOC, dividing by $\alpha_m W_{jt}^{\varrho-1} s_{jt}$, and defining $\kappa_m^{(\varrho)} \equiv \sum_{f\in\mathcal{J}_{mt}} (1/\Delta_{ft} - 1) W_{ft} (1 - s_{0t|m}) s_{ft|m}$, we have,

$$\left(\frac{1}{\Delta_{jt}^{(\varrho),BertColl}} - 1 \right) \left[1 - (1 - s_{0t|m}) s_{jt|m} (1-\lambda_m) \right] = \frac{1}{\alpha_m W_{jt}^{\varrho}} + \frac{\lambda_m \kappa_m^{(\varrho)}}{W_{jt}}.$$

As in the log-wage case, this is a system of J_m equations in J_m unknowns and requires numerical solution.

Identifying ϱ. The wage transform adds a single parameter. Because ϱ enters the within-market labor supply equation nonlinearly, it is not eliminated by the differencing that identifies α_m in Section 4.2: differencing across firms j and f in the same market gives

$$\log L_{ft} - \log L_{jt} = \frac{\alpha_m}{\varrho} \left(W_{ft}^{\varrho} - W_{jt}^{\varrho} \right) + (\xi_{ft} - \xi_{jt}),$$

so (α_m, ϱ) are jointly identified from the within-market wage variation, estimated by nonlinear IV using the labor demand instrument Z_{ft}^{ξ} of Section 5.3. For example, the two parameters can be recovered from the moments

$$\mathbb{E}\left[\left\{ (\log L_{ft} - \log L_{jt}) - \frac{\alpha_m}{\varrho} \left(W_{ft}^{\varrho} - W_{jt}^{\varrho} \right) \right\} Z_{ft}^{\xi} \right] = 0$$

and

$$\mathbb{E}\left[\left\{ (\log L_{ft} - \log L_{jt}) - \frac{\alpha_m}{\varrho} \left(W_{ft}^{\varrho} - W_{jt}^{\varrho} \right) \right\} \left(Z_{ft}^{\xi} \right)^2 \right] = 0,$$

similar to Birchall et al. (2024). Because the residual is nonlinear in ϱ, these two moments are non-redundant even though they are constructed from a single instrument: a continuous Z^{ξ}_{ft} traces out the curvature of the wage response that separates ϱ from α_m, so no second instrument is required. Setting $\varrho \to 0$ collapses the residual to its log-case counterpart and the first moment reduces to equation (13).

Identifying s_0. Identification of the outside share $s_{0t|m}$ deviates substantively under Box-Cox utility relative to Appendix C: a uniform shift Δ to all log wages in the market no longer rescales the wage index by the common factor $e^{\alpha_m \Delta}$. Under the shift, firm f's wage becomes $W_{ft}e^{\Delta}$, so

$$\delta_{ft}(\Delta) = \alpha_m \frac{(W_{ft}e^{\Delta})^{\varrho} - 1}{\varrho} + \xi_{ft}, \qquad \frac{\partial \delta_{ft}}{\partial \Delta} = \alpha_m (W_{ft}e^{\Delta})^{\varrho}.$$

Then,

$$\left.\frac{\partial s_{0t|m}(\Delta)}{\partial \Delta}\right|_{\Delta=0} = -s^2_{0t|m} \sum_{f \in \mathcal{J}_{mt}} \exp(\delta_{ft})\, \alpha_m W^{\varrho}_{ft} = -s_{0t|m} \sum_{f \in \mathcal{J}_{mt}} s_{ft}\, \alpha_m W^{\varrho}_{ft}.$$

Since $\log \bar{L}_{mt} = \log(1 - s_{0t|m}) + \log \bar{L}^{\text{pot}}_{mt}$ and $s_{ft} = (1 - s_{0t|m}) s_{ft|m}$, the factor $(1 - s_{0t|m})$ cancels and

$$\left.\frac{\partial \log \bar{L}_{mt}}{\partial \Delta}\right|_{\Delta=0} = \frac{s_{0t|m} \sum_{f \in \mathcal{J}_{mt}} s_{ft}\, \alpha_m W^{\varrho}_{ft}}{1 - s_{0t|m}} = s_{0t|m} \sum_{f \in \mathcal{J}_{mt}} s_{ft|m}\, \alpha_m W^{\varrho}_{ft}.$$

The aggregate employment response therefore depends on the within-market distribution of W^{ϱ}_{ft} rather than collapsing to the scalar $\alpha_m s_{0t|m}$.

Instrumenting the market-average wage with the aggregate demand shock ζ_{mt} as in Appendix C,

$$s_{0t|m} \Xi^{(\rho)}_{mt} = \frac{\text{Cov}(\log \bar{L}_{mt},\, \zeta_{mt})}{\text{Cov}(\bar{w}_{mt},\, \zeta_{mt})}, \quad \text{where} \quad \Xi^{(\rho)}_{mt} \equiv \sum_{f \in \mathcal{J}_{mt}} s_{ft|m}\, \alpha_m W^{\varrho}_{ft}.$$

The multiplier $\Xi^{(\rho)}_{mt}$ can be constructed using only the α_m, ϱ estimates and observables, so $s_{0t|m}$ is identified once α_m, ϱ are identified.

Conduct testing. Given α_m, ϱ, $s_{0t|m}$, each firm's markdown can be constructed for each candidate conduct, and the RV test statistic implementation then proceeds exactly as in the main text.