ROCKWOOL FOUNDATION BERLIN
INSTITUTE FOR THE ECONOMY AND THE FUTURE OF WORK

CReAM Centre for Research & Analysis of Migration

DISCUSSION PAPER SERIES 31/24

Monopsony Power in the Labor Market

José Azar, Ioana Marinescu

DECEMBER 2024

ROCKWOOL Foundation Berlin
www.rfberlin.com

Centre for Research & Analysis of Migration
www.cream-migration.org

Monopsony Power in the Labor Market

José Azar and Ioana Marinescu

June 2024

Abstract

Labor economics often assumes that wages w are equal to the marginal revenue product of labor $MRPL$. However, recent literature has shown that firms' market power allows them to pay wages substantially below marginal productivity. The markdown $(MRPL - w)/w$ is our preferred measure of firms' monopsony power, and captures the percent wage increase that would occur if monopsony power were eliminated. We derive the markdown across three classes of models, each embodying a distinct source of monopsony power. First, in oligopsony models, monopsony power arises from strategic interactions between large firms, and is related to labor market concentration. Second, in job differentiation models, monopsony power arises from workers' heterogeneous preferences over jobs that differ in wages and amenities. Finally, in search and matching models, it arises from frictions that prevent workers from accessing all existing job vacancies. To identify the markdown, empirical studies often rely on estimating the firm-level labor supply elasticity and taking its inverse as a measure of the markdown. A few studies directly estimate $MRPL$ using a production function approach. Across studies, the markdown typically ranges between 15% and 50% implying that wages would increase by 15 to 50% if firms' monopsony power were eliminated. Finally, we analyze the policy implications of monopsony power in three areas, drawing on both theory and empirical analysis: merger control in antitrust policy, the regulation of non-competition agreements, and minimum wages. Monopsony power helps explain how mergers and non-competition agreements can lower wages, and how minimum wages can increase employment. Overall, the literature shows that monopsony power is significant, and should be considered when analyzing policy and the sources of wage variation.

Keywords: monopsony, oligopsony, markdown, wages, labor market concentration, labor supply elasticity, antitrust, mergers, imperfect competition, minimum wage

Word count:19,804.

1 Introduction

While the labor market has traditionally been studied under the assumption of perfect competition, a recent theoretical and empirical literature,[1] – reviewed in Sokolova and Sorensen (2021); Manning (2021); Card (2022) – has challenged this assumption by documenting the significant monopsony

[1] Manning (2011) reviewed the older literature relevant to imperfect competition in the labor market in a previous volume of this Handbook.

power that firms wield. For the purposes of this chapter, we focus on the markdown as the key measure of monopsony power, i.e. firms' market power as buyers of labor services.[2] The markdown is the gap between workers' marginal revenue product ($MRPL$) and the wage[3] w (scaled by the wage), i.e. $(MRPL-w)/w = MRPL/w-1$.[4] The markdown measures the percent wage increase that would occur, all other things equal, if employers' monopsony power were eliminated so that the wage is equal to the marginal revenue product of labor, $w = MRPL$.

This chapter summarizes and explains the insights from this burgeoning monopsony literature by focusing on three pivotal theoretical frameworks: oligopsony, job differentiation, and search and matching. These frameworks not only deepen our understanding of the sources of monopsony power but also have implications for public policy, particularly in the areas of merger control, non-competition agreements, and minimum wage regulation.

In the oligopsony framework, we explain how to extend the basic "isolated firm" monopsony model to introduce strategic interactions between firms. Firms compete by selecting employment levels, taking the market-level labor supply curve as given. In this framework, the average markdown in a labor market increases with concentration, and decreases with the labor supply elasticity to the market.

The job differentiation framework posits that monopsony power can arise even in markets with atomistic firms, due to heterogeneous worker preferences over jobs that provide both wages and amenities. The greater the degree of job differentiation, the greater the markdown. This framework is instrumental in understanding how firms can pay wages below the marginal revenue product of labor without losing all of their workers. Workers are willing to accept less than competitive wages when the job provides higher utility – a better bundle of wages and amenities – than the next best alternative. Job differentiation can serve as a microfoundation for the finite labor supply elasticity to the market in the oligopsony model. New theoretical frameworks have added large firms, allowing for the integration of oligopsony and job differentiation. In these hybrid models, both job differentiation and labor market concentration increase the markdown.

The search and matching framework introduces search frictions. Because of these frictions, workers cannot instantaneously meet all available jobs, and thus cannot fully benefit from competition between employers. These frictions enable even atomistic firms to offer wages below what would prevail in a perfectly competitive market. The recent integration of large firms into these models offers fresh perspectives on wage-setting dynamics and introduces labor market concentration as a negative determinant of wages. In search and matching models, the markdown typically decreases with the worker's non-employment income (the outside option), and varies with concentration and the nature and degree of search frictions.

Each of these theoretical frameworks has been used to quantify the role of monopsony power in the labor market using constructs like labor market concentration, the labor supply elasticity,

[2]As we discuss below, "competition" and "market power" have complex meanings that cannot be fully captured by the markdown.

[3]In a broader sense, the wage can be understood as full compensation, including costly amenities provided by the firm, as we will discuss.

[4]More generally, one may define monopsony power as the ability of employers to pay workers less than the *competitive* wage, while recognizing that the competitive wage benchmark may be different from the marginal revenue product, depending on model assumptions. For example, if the solution to the firms' profit maximization problem is not interior, then their level of employment may not satisfy the first-order condition that the wage equals the value of the marginal product of labor.

and the production function. We summarize the results of this empirical literature and the econometric challenges that it seeks to address. Starting with concentration, the empirical research often focuses on wages as the outcome instead of markdowns. Yet, wages need not decrease with concentration, even if theory predicts that the *markdown* increases with concentration. Wages might increase with concentration if the dominant firm's productivity significantly improves, allowing it to expand its market share. Thus, concentration is a useful proxy for market power (i.e., the markdown), but it may not be as informative about the wage, and ultimately worker *welfare*. Nevertheless, the empirical literature typically finds that higher concentration is associated with lower wages. We then summarize results on the labor supply elasticity and the markdown (see Tables 4, 5 and 6). The literature typically finds a labor supply elasticity to the firm between 2 and 6, a labor supply elasticity to the market between 0.5 and 5, and a markdown between 15% and 50%. The markdown estimates thus imply that, all other things equal, wages would increase between 15% and 50% if monopsony power could be eliminated. Overall, the literature paints a consistent picture indicating that firms have significant monopsony power in the labor market.

Monopsony power affects our understanding and interpretation of variation in wages across workers, firms, and over time. In general, we should be careful not to interpret the wage as a direct measure of the marginal productivity of labor – a common assumption in economic analysis. Variation in wages may be driven by variation in the markdown as well as variation in productivity. Future research should consider whether the role of monopsony power is significant to understand the economic question at hand.

The models of monopsony power we discuss in this chapter can fundamentally affect our understanding of policy effects in various contexts. We focus on three areas where there exists a significant body of empirical research. First, the oligopsony model elucidates the mechanisms that drive the observed wage suppression in mergers that increase labor market concentration. Second, the search and matching model can inform the policy debate around non-competition agreements, and helps explain the observed negative effects on wages and labor mobility. Lastly, the oligopsony and job differentiation models offer a lens to examine the disparate impacts of minimum wage policies across labor markets with differing levels of monopsony power. Empirical studies of minimum wage effects suggest that markdowns can buffer the labor market against the negative employment effects predicted by the perfectly competitive model; the minimum wage can even *increase* employment in the least competitive labor markets. Monopsony power thus helps explain why many studies fail to find a negative effect of the minimum wage on employment.

The chapter proceeds as follows. In Section 2, we describe and compare the three theoretical frameworks that help explain monopsony power: oligopsony, job differentiation, and search and matching. We provide formulas for the markdown (Table 1) and the firm-level and market-level labor supply elasticities (Table 2) in these different models. In Section 3, we discuss the empirical challenges in estimating monopsony power, and summarize the results from various strands of the literature, including the labor supply elasticity, labor market concentration, structural estimation, and the production function approach. In Section 4, we discuss the role of monopsony power in policy analysis, focusing on areas where the literature is most developed and most well integrated with theory: merger control, non-competition agreements, and the minimum wage.

2 New quantitative models of monopsony power

2.1 What is monopsony power?

The markdown – the difference between workers' marginal revenue product and the wage – is in our view the key measure of monopsony power. A primary goal of this theory section is to derive formulas for the markdown and its determinants across models, so that we can better understand what drives market power in different models.

Before we get started with our theoretical discussion, it is important to recognize that "market power" and "competition" have a rich set of meanings in policy and business contexts. For example, in 2023, the White House Office of Budget and Management (OIRA, 2023) defined these terms as follows:

> *Competition* is, among other things, the process by which individuals or firms vie to win customers' business for goods or services, to purchase suppliers' goods or services, or to hire workers for their labor services. Competitive markets are characterized by (but not exclusively characterized by) the presence of independent and rival buyers and sellers such that each market participant has many potential options to turn to. Encouraging competitive markets is an important policy goal. Competitive markets are associated with lower prices for consumers, higher wages for workers, more innovative products and services, more business formation, and greater resilience to unexpected events. When markets are less competitive, we say that certain market participants have *market power*. Firms with market power have the ability to change their behavior so as to increase their own profits or advance their other interests at the expense of others.

If we adopt a broad definition of competition and market power (such as the one above), the markdown does not capture all the nuances of market power. The markdown does not directly measure how the competitive process unfolds in a particular labor market. Instead, we can think of the markdown as a measure of the outcome of this competitive process, which gives us information about firms' ability and willingness to reduce wages below the marginal revenue product of labor. Our theoretical discussion helps make more precise what mechanisms drive the markdown in specific models. Ultimately, the markdown is a useful summary statistic for market power in the labor market at a given point in time.

2.2 Oligopsony

Oligopsony is the classical theoretical framework where firm size matters, and explains how wage determination is affected by monopsony power. Labor market concentration is a key determinant of the markdown in this class of models.

In the textbook perfectly competitive labor market, infinitely many firms compete for workers, and wages are equal to the worker's marginal revenue product. These firms are assumed to be atomistic, so that they cannot affect market outcomes. At the other extreme of the perfectly competitive model is the literal case of monopsony, in which there is a single employer in a labor market. The simplest model of monopsony, the "isolated firm" model, goes back to Robinson

([1933] 1966). In contrast to the perfectly competitive labor market model where the individual firm faces a constant labor supply curve, this model considers an upward sloping labor supply curve.

Oligopsony models extend the basic "isolated firm" monopsony model to allow for multiple large firms operating in a labor market. The term "large" indicates that these firms are significant in size relative to the market; they are not atomistic or infinitesimal. Consequently, these large firms influence market-level outcomes, such as employment levels. This scenario introduces strategic interactions among firms in the labor market. Coupled with the upward-sloping market-level labor supply curve, these strategic interactions form the foundational elements of the oligopsony model. These components differentiate the oligopsony model from a model of perfect competition. The oligopsony model is sometimes dubbed "classical monopsony" in the literature.

In the simplest version of the oligopsony model,[5] several large firms compete with each other by selecting the level of employment, given their expectations of competing firms' employment levels. The first-order condition for a firm n in the oligopsonistic labor market is:

$$R'(L_n) = w(L) + w'(L)L_n, \tag{1}$$

where $R'(L_n)$ is the marginal revenue product of labor, L is market-level employment, $w(L)$ is the inverse labor supply curve to the market, and L_n is the employment of firm n. In the perfectly competitive model, the employment level of the firm would not influence the equilibrium wage in the market, so the term $w'(L)L_n$ would be set to zero. As a result, we would have a wage that is equal to the marginal revenue product of labor, a well-known result for perfectly competitive markets. Therefore, maximum competition in this model results in wages equal to marginal productivity.

It's worth noting here that the firm is assumed to pay the same wage to all its workers, and this assumption leads to the term $w'(L)L_n$ in equation (1). This equality constraint is a fundamental assumption that leads to lower employment: because the firm has to pay incumbent workers more in order to hire the marginal worker, it chooses a lower level of employment in equilibrium.

A firm's markdown (the difference between the marginal revenue product and the wage as a share of the wage) is given by:

$$\frac{R'(L_n) - w}{w} = \frac{s_n}{\eta}, \tag{2}$$

where $R'(L_n)$ is the revenue marginal product of labor of firm n, w is the equilibrium wage in the labor market, s_n is the employment share of firm n, and η is the market-level elasticity of labor supply. If the firm is a literal monopsonist, the firm's share is equal to one and the market-level elasticity of labor supply is also the firm's elasticity of labor supply. In an oligopsony labor market, each firm's labor supply elasticity is larger than the monopsonist's, since each firm's share is less than one, and firms compete with each other.

The average markdown or Pigou's rate of exploitation is equal to the market-level inverse elasticity of labor supply times the employment Herfindahl-Hirschman Index (HHI) in the labor market (Boal and Ransom, 1997):

$$\sum_{n=1}^{N} s_n \frac{R'(L_n) - w}{w} = \frac{\sum_{n=1}^{N} s_n^2}{\eta} = \frac{HHI}{\eta}, \tag{3}$$

[5] For example, Azar, Huet-Vaughn, Marinescu, Taska, and von Wachter (2023) use this oligopsony model and derive its wage and employment effects.

where the HHI is the defined as the sum of the squares of employment shares for each firm: $HHI = \sum_{n=1}^{N} s_n^2$.

An alternative model of oligopsony would feature firms competing in wages (i.e., analogous to competing à la Bertrand in product markets) as opposed to firms competing in levels of employment (i.e., the analogue of competing à la Cournot in product markets). With undifferentiated jobs and symmetric firms, competition in wages leads to a "Bertrand paradox" in which the Nash equilibrium wage is equal to the marginal product of labor as long as there are two or more firms competing in the labor market (Bertrand, 1883). The reason is that, as long as there is a gap between the marginal revenue product of labor and the wage, any firm in the market has an incentive to offer a wage infinitesimally above the prevailing wage, and hire all the workers in the labor market. This implies that the only equilibrium wage is equal to the marginal product of labor, and thus the wage outcome is the same as under perfect competition. This model is generally considered unrealistic for modeling oligopsony in labor markets but has some application in models with search frictions, which we discuss later. A model of firms competing in wages is more appealing when jobs are differentiated, as job differentiation can allow for equilibrium wages that are below the marginal revenue product of labor.

2.3 Differentiated jobs

In the Cournot oligopsony model, there is a finite labor supply elasticity to the whole labor market (as opposed to the individual firm). One way to justify this assumption is to posit that jobs in a labor market are differentiated from jobs in other labor markets. Therefore, at least implicitly, the homogeneous jobs Cournot model already hints at the idea of differentiation as a source of monopsony power. However, one can also add differentiation explicitly both across markets and across jobs in the same labor market.

In a differentiated jobs model, the key source of monopsony power is not the finite number of firms as in oligopsony models, but rather the fact that workers have heterogeneous preferences over jobs that differ in wages and amenities. When a worker's job is better (provides higher utility) than their next best option, the firm is able to pay the worker less than the marginal revenue product of labor.

When jobs are differentiated and workers' preferences are heterogeneous, a firm that marginally increases its wage above competitors' levels will not attract all employees from the market. Since workers weigh wages against other job amenities differently, only a negligible number of them will transition to the firm that decides to offer a marginally higher wage. Consequently, even amidst wage-setting competition among firms, equilibrium wages can remain below the marginal product of labor. Introducing job differentiation can therefore resolve the Bertrand paradox.

The models we explore in this section are typically premised on the notion that amenities are exogeneously provided by firms at no cost, sidestepping the specifics of how amenities are determined. While the seminal Rosen compensating differentials model (Rosen, 1974, 1986) accounts for the cost of amenities, its integration with monopsony power is not straightforward and remains a subject of active inquiry (Lavetti, 2023). The Rosen framework assumes that amenities are a cost that negatively impacts productivity, and that firms optimally choose the level of amenities depending on the marginal cost of providing them. It also further assumes that wages are equal to marginal productivity. Assuming that wages are competitive is important in this literature to draw inferences about the amenities of different jobs from data on wage differentials (Lavetti, 2023). However, with monopsony power, wages are not equal to marginal productivity. This makes it difficult to use Rosen-style models to estimate the value of amenities in the presence of monopsony power.[6]

Here, we delve into the details of differentiated jobs models with exogenous amenities, which can more easily accommodate monopsony power. These models fall into two categories: representative agent models (Bhaskar, Manning, and To, 2002; Berger, Herkenhoff, and Mongey, 2022a) and discrete choice models of labor supply (Card, Cardoso, Heining, and Kline, 2018; Azar, Berry, and Marinescu, 2019). While discrete choice models are more widely used in empirical work, we start with a discussion of the representative choice model for several reasons. First, it is a more straightforward theoretical extension of the oligopsony model, as it assumes Cournot competition, while discrete choice models typically assume Bertrand competition. Additionally, the representative agent model with Cournot competition is more tractable, resulting in a simple formula for the markdown that includes a parameter for job differentiation. Within the representative choice model, we will first assume atomistic firms so that markdowns arise from job differentiation alone rather than from a finite number of firms. We then assume a finite number of firms, enriching the job differentiation model with oligopsony interactions.

2.3.1 Monopsonistic Competition with a Representative Household and Differentiated Jobs

In the representative household model of Berger, Herkenhoff, and Mongey (2022a), there is one household with a utility function that represents preferences for differentiated jobs, and it supplies some of its labor to each of the various jobs in the economy. This is analogous to a consumer choice model in which the representative household chooses the share of its budget that it will allocate to each good. If a job becomes more attractive, for example due to a higher wage, the representative household increases the share of its labor endowment that it supplies to that job. Again, as in the consumer choice model, unless the choices are perfect substitutes (in this case, perfectly substitutable jobs), a small change in the price (i.e., wage) will not induce a discrete jump in the market share of that choice. Thus, job differentiation induces an upward sloping labor supply to the firm.[7]

Consider a representative household that has preferences over consumption C and a bundle of labor supply to differentiated jobs L, and l_i is the labor supply to firm i. The utility U of the representative household is:

$$U(C, L) = C - \frac{1}{\overline{\varphi}^{\frac{1}{\varphi}}} \frac{L^{1+\frac{1}{\varphi}}}{1 + \frac{1}{\varphi}}, \tag{4}$$

[6] Dube, Naidu, and Reich (2022) illustrate the limits of the Rosen model to understand amenity provision. They show that, when firms have monopsony power, minimum wages need not decrease the provision of costly amenities, contrary to this model's predictions. Empirically, they find no effect of Walmart's company-imposed minimum wage on amenities provided by the firm.

[7] Common specifications for the representative agent model include quadratic utility, which leads to a linear system of labor supplies, and constant-elasticity of substitution, which leads to a constant elasticity of labor supply.

where

$$L = \left[\int_0^1 l_i^{\frac{\theta+1}{\theta}} di \right]^{\frac{\theta}{\theta+1}} \tag{5}$$

This preference specification generates constant-elasticity market-level and firm-level labor supply functions. The φ parameter determines the aggregate wage elasticity of labor supply, while θ determines the level of job differentiation in the market, which affects the firm-level elasticity of labor supply ($\overline{\varphi}$ is a labor supply shifter). When θ goes to infinity, jobs become perfect substitutes, and the model becomes the same as a Cournot oligopsony with homogeneous labor (Section 2.2). The household's budget constraint is $pC = \int_0^1 w_i l_i di + b$, where b is nonlabor income.

The first-order condition with respect to l_i yields the inverse labor supply function to firm i (which takes aggregate labor supply L as given):

$$w_i = \frac{1}{\overline{\varphi}^{\frac{1}{\varphi}}} L^{\frac{1}{\varphi}} \left(\frac{l_i}{L} \right)^{\frac{1}{\theta}} \tag{6}$$

We can rewrite this in terms of the wage index $W = \left[\int_0^1 w_j^{1+\theta} dj \right]^{\frac{1}{1+\theta}}$ as follows:

$$w_i = \left(\frac{l_i}{L} \right)^{\frac{1}{\theta}} W \tag{7}$$

In the equation above, the wage of the individual firm depends on its market share l_i/L and the level of job differentiation θ. However, the wage does not depend on the elasticity of labor supply at the market level φ, because the firm does not take into account its effect on aggregate employment L. In the next subsection, with oligopsonistic competition, firms *will* take into account their effect on aggregate employment L.

Firm i chooses l_i (or, equivalently, w_i) to maximize profits π_i:

$$\pi_i = pF(l_i) - w_i(l_i) l_i \tag{8}$$

The firm takes as given the price of its product (there is therefore an assumption of perfect competition in the product market), as well as the aggregate wage and labor supply. However, while the firm cannot affect the overall level of wages, it has monopsony power over its own wage, because the jobs it offers are differentiated from other firms' jobs. As in the classical monopsony model, the first-order condition with respect to l_i implies that firms equate the marginal revenue product of labor to the marginal cost of labor, with the latter being above the wage:

$$pF'(l_i) = w_i \left(1 + \frac{1}{\theta} \right) \tag{9}$$

Note that in this case, $pF'(l_i)$ is the value of the marginal product of labor and also equal to the marginal revenue product of labor $R'(l_i)$ (see equation 1). As we will discuss in section 2.3.6 below (see specifically equation (40)), when there is product market power in addition to labor market power, the value of the marginal product of labor is different from the marginal revenue product of labor.

The markdown is given by:

$$\mu_i \equiv \frac{pF'(l_i) - w_i}{w_i} = \frac{1}{\theta} \tag{10}$$

Thus, the markdown is positive even if, in this case, the firms are atomistic, and the level of labor market concentration is equal to zero because there is an infinite number of firms. As θ goes to infinity, jobs become homogeneous, and the markdown goes to zero because there is no longer any job differentiation. Further, the markdown does not depend on the labor supply elasticity to the market φ because, as we noted above, the firm does not take into account its effect on aggregate employment.

2.3.2 Oligopsony with a Representative Household and Differentiated Jobs

If the number of firms is finite, then we can write a similar model but with firms having strategic interactions, and the markdown will depend on the HHI. In this case, firms do not take as given the aggregate wage and aggregate employment. If there are N firms, the labor supply to a firm i is:

$$w_i = \left(\frac{l_i}{L}\right)^{\frac{1}{\theta}} W, \tag{11}$$

where $L = \left[\sum_{j=1}^{N} l_j^{\frac{\theta+1}{\theta}} dj\right]^{\frac{\theta}{\theta+1}}$ and $W = \left[\sum_{j=1}^{N} w_j^{1+\theta} dj\right]^{\frac{1}{1+\theta}}$.

The first-order condition of firm i again equalizes the marginal revenue product of labor to the marginal cost of labor. However, because firms now perceive the effect of their actions over aggregate employment and wages, the slope of the inverse labor supply with respect to labor now has an extra term that depends on firm i's payroll market share[8] $\tilde{s}_i = w_i l_i/(WL)$. The first order condition becomes:

$$pF'(l_i) = w_i \left(1 + \frac{1}{\theta} + \left(\frac{1}{\varphi} - \frac{1}{\theta}\right) \tilde{s}_i\right) \tag{12}$$

This implies that the markdown of firm i is given by:

$$\mu_i \equiv \frac{pF'(l_i) - w_i}{w_i} = \frac{1}{\theta} + \left(\frac{1}{\varphi} - \frac{1}{\theta}\right) \tilde{s}_i \tag{13}$$

Taking a weighted average, weighted by payroll shares, yields an expression in terms of the (payroll-share) Hefindahl-Hirschman Index:

$$\sum_{i=1}^{N} \tilde{s}_i \mu_i = \frac{1}{\theta} + \left(\frac{1}{\varphi} - \frac{1}{\theta}\right) HHI \tag{14}$$

Relative to the markdown without job differentiation (see equation 3), there are two main differences: (i) the $HHI = \sum_{i=1}^{N} \tilde{s}_i^2$ is in terms of payroll shares instead of employment shares s_i, and (ii) there is a constant term $1/\theta$ that reflects the fact that, with differentiation, firms have a baseline level of market power even if concentration is zero (i.e., even if they are atomistic). The fact that the HHI is based on payroll instead of employment shares only matters if the firms do not pay the same wage (e.g. because they have different levels of productivity as in Berger, Herkenhoff, and Mongey (2022a)).

[8] If firms are identical, the payroll market shares are the same as the employment shares and equal to $1/N$.

2.3.3 Monopsonistic Competition and Oligopsony with Discrete Choice

In contrast to representative household models, workers in a *discrete* choice model only pick one job rather than spreading themselves across jobs. Discrete choice models can be used to analyze how workers make decisions about job choice. These models assume that workers pick the job that gives them the highest utility. Here again, the markdown will depend on firm shares, and on the HHI.

Consider a labor market with J firms offering differentiated jobs. Firms compete by setting wages (Bertrand) and forming expectations about the wages of their competitors. There is a continuum of workers of measure L that have random utility over the jobs. In particular, worker i's utility for firm j's job is:

$$u_{ij} = \delta_j + \epsilon_{ij}, \tag{15}$$

where δ_j is the job's mean utility, and ϵ_{ij} is the random component of utility, which has an extreme value distribution. δ_j allows for jobs to be different in a deterministic way, and these differences are valued in the same way by all workers. This can be labeled as vertical differentiation. The random component of utility adds more job differentiation into the model, as a given worker values two jobs differently even if they have the same δ_j. This can be labeled as horizontal differentiation.[9] Because ϵ_{ij} is random across jobs *and* workers, the ranking of jobs is different across workers. The mean utility is $\delta_j = \alpha + \eta \log(w_j)$. Because the wage is a component of δ_j, two jobs with the same wage can be valued differently by the same worker due to the random component ϵ_{ij}. With these assumptions, it can be shown that the labor market share of firm j is:

$$s_j = \frac{e^{\delta_j}}{\sum_{k=1}^{J} e^{\delta_k}} \tag{16}$$

This implies that employment at firm j is equal to $s_j \times L$, so we can rewrite this as:

$$\frac{l_j}{L} = \frac{w_j^\eta}{\sum_k w_k^\eta} \tag{17}$$

Firm j chooses its wage to maximize its profits:

$$\max_{w_j} (A_j - w_j)s_j, \tag{18}$$

where A_j is the marginal revenue product of labor.

The first-order condition of firm j is:

$$-s_j + (A_j - w_j)\frac{\partial s_j}{\partial w_j} = 0, \tag{19}$$

where

$$\frac{\partial s_j}{\partial w_j} = \frac{\eta s_j (1 - s_j)}{w_j} \tag{20}$$

[9] For example, Lamadon, Mogstad, and Setzler (2019) use this terminology.

From the first-order condition, we obtain the formula that equates the markdown and the inverse-elasticity of labor supply, which in this case is:

$$\frac{A_j - w_j}{w_j} = \frac{s_j}{w_j}\frac{1}{\partial s_j/\partial w_j} = \frac{1}{\eta(1-s_j)} \tag{21}$$

This formula is similar to equation (13) in the prior section 2.3.2, with the labor supply elasticity parameters (respectively φ above, and η here) decreasing the markdown while the shares (payroll share $\tilde{s}_i$ above, and share s_j here) increase the markdown. However, the formulas are not identical because here we have Bertrand rather than Cournot competition.

In the prior section 2.3.2, there was an explicit job differentiation parameter θ. The reader may be wondering how job differentiation is parameterized in this model. Imagine there was no random component of utility ϵ_{ij} in equation (15). In that case, all jobs would be exactly the same, except maybe for the wage. Thus, if a firm reduces the wage below the maximum that is paid by other firms in the labor market, it would get zero workers. Therefore, firms would have no market power without the random component. The fact that the random component has a positive variance is what introduces worker-specific job differentiation, and through it a finite labor supply elasticity to the firm.[10]

If firms are symmetric in the sense that $A_j = A$ for all j (and the deterministic component of utility is also equal across firms except maybe for the wage, i.e., $\delta_j = \alpha + \eta \log(w_j)$ as we have assumed so far), there is an exact relation between the markdown and the HHI:

$$\frac{A - w}{w} = \frac{1}{\eta(1-HHI)}, \tag{22}$$

where $HHI = \frac{1}{J}$, and there is a closed-form solution for the wage as well:

$$w = \frac{A}{1 + \frac{1}{\eta(1-HHI)}} \tag{23}$$

If firms are not symmetric so that each firm has its own productivity A_j, the markdown in equation (21) doesn't imply an exact relationship between the average markdown and market concentration. However, we can derive an approximate relationship if we assume that firms have similar but not identical market shares.[11]

[10] If there were a variance parameter for the random component ϵ_{ij}, this parameter would play a role similar to the θ in the model in the previous section. However, it is convention to fix the variance of the random component in discrete choice models, because, econometrically, this variance cannot be identified separately from the deterministic component parameters (i.e. the coefficients on the observed variables that predict worker choice, such as the wage); indeed, what really matters is the relative magnitude of the random component and the deterministic component. Therefore, with a fixed variance of the random component, it is the magnitude of the deterministic component parameters (relative to that fixed variance) that drives differentiation in this model.

[11] The expression using a second-order Taylor approximation around $\frac{1}{J}$ is:

$$\sum_j s_j \frac{A_j - w_j}{w_j} \approx \frac{1}{\eta}\left[\frac{J(J^2-3J+1)}{(J-1)^3} + \frac{J^3}{(J-1)^3}HHI\right] \tag{24}$$

Under some assumptions about hours worked (in particular, if workers target a given level of income), it can be shown that the logit model in this section is equivalent to the CES model developed in the previous section 2.3.2 (Berger, Herkenhoff, and Mongey, 2019).[12] If instead of Bertrand competition in wages, we assumed that firms compete in employment (Cournot), then the logit model would imply the same relationship between the average markdown and the payroll share HHI as in the representative household CES model, because the logit and CES models are equivalent in the sense that they imply the same labor supply system.

2.3.4 Nested logit

One can build more flexible discrete choice models by using nesting.[13] In a nested model, the worker first makes a decision whether to work in a given labor market or not; if the worker decides not to work in that given labor market, the worker chooses the "outside option". The outside option represents either nonemployment or working in other labor markets. If the worker chooses to work in the market, they next choose which firm to work for within the market. Thus, we have a nested logit framework with two nests: a top-level nest for the decision of whether to work in the market or not, and a bottom-level nest for the decision of which firm in the market to work for. Azar, Berry, and Marinescu (2022) estimate such a nested model.

Formally, the utility of worker i for working for firm j is:

$$u_{ij} = \alpha + \eta \log(w_j) + v_i(\lambda) + \lambda\epsilon_{ij} \tag{25}$$

Here, $\delta_j \equiv \alpha + \eta \log(w_j)$ is the deterministic component of utility. There are two error terms in the utility function: an extreme value error term ϵ_{ij}, which captures random shocks to the utility of worker i from working at firm j, and a group preference term $v_i(\lambda)$, which represents the worker's overall preference for the group, i.e. the labor market they are considering. This group term has a unique distribution that ensures that the term $v_i(\lambda) + \lambda\epsilon_{ij}$ is also distributed as an extreme value. The nesting parameter $\lambda \in [0, 1]$ determines how differentiated the jobs inside the market are from the outside option. When $\lambda = 1$, the model simplifies to a standard logit model, where the outside option is treated as part of the same nest as the focal labor market. This means that jobs within the focal labor market compete just as much with jobs outside the market as within the market, implying no differentiation between the two markets in terms of the worker's choice. Conversely, when $\lambda = 0$, then the labor market is fully segmented from the outside option, and a worker either always works in the market, or always chooses the outside option.

The labor market share of firm j in the nested logit model is:

$$s_j = \underbrace{s_{j|g}}_{\text{within-group share}} \times \underbrace{s_g}_{\text{group share}} = \frac{e^{\frac{\delta_j}{\lambda}}}{e^I} \times \frac{e^{\lambda I}}{1 + e^{\lambda I}}, \tag{26}$$

[12] Berger, Herkenhoff, and Mongey (2019) show this is true in the nested logit and nested CES case. Since the non-nested versions are special cases of the nested imposing specific parameter values, their proof also works for the non-nested case. The analogous arguments for the product market logit and CES and nested logit and nested CES are made by Anderson, De Palma, and Thisse (1988) and Verboven (1996), respectively. Specifically, what these papers show is that a "fictitious" representative consumer with (nested) CES preferences can represent a population of consumers with heterogeneous (nested) logit preferences.

[13] Similarly, nesting can be used to obtain more flexible versions of the representative household CES model.

where

$$I = \log \sum_{k=1}^{J} e^{\frac{\delta_k}{\lambda}} \tag{27}$$

is the "inclusive value", which measures the expected maximum utility of working in the specific labor market group. The group share is the share of workers that work at any job in this particular labor market. The within-group share is the share of firm j within the labor market g.

Assuming that firms compete in wages, the first-order condition of firm j is the same as in the simple logit case, except that the formula for the slope of the market share with respect to the wage now also depends on the nesting parameter and is given by:

$$\frac{\partial s_j}{\partial w_j} = \frac{\eta}{w_j} \frac{1}{\lambda} s_j \left[1 - (1-\lambda) s_{j|g} - \lambda s_j\right] \tag{28}$$

The own-wage elasticity of labor supply to the firm is:

$$\frac{\partial \log s_j}{\partial \log w_j} = \frac{\eta}{\lambda} \left[1 - (1-\lambda) s_{j|g} - \lambda s_j\right] \tag{29}$$

The markdown of firm j is the inverse of the elasticity right above:

$$\frac{A_j - w_j}{w_j} = \frac{\lambda/\eta}{1 - (1-\lambda) s_{j|g} - \lambda s_j} \tag{30}$$

There is no closed-form solution for the equilibrium markdown, given that the market shares depend on the wage, and that, even in the symmetric case, market shares are not equal to $1/J$ because of the outside option. However, empirically, if we know the elasticity and the market shares, we can calculate the markdown without numerically solving the nonlinear system of first-order conditions.

2.3.5 Nesting with monopsonistic competition

Lamadon, Mogstad, and Setzler (2022) develop an equilibrium model of the labor market with two-sided heterogeneity where workers view firms as imperfect substitutes because of heterogeneous preferences over non-wage job characteristics, i.e. amenities. Such heterogeneous preferences lead to a finite labor supply elasticity to the firm. They define markets as an industry by commuting zone. They allow for correlation of worker idiosyncratic preferences for firms within each nest, with the degree of correlation varying across nests. Each nest is assumed to include many firms, and firms therefore are assumed not to act strategically, that is, firms do not take into account the impact of changing their own wages on the market-level wage. Still, firms do exercise market power by taking into account the finite labor supply elasticity to the firm.

The assumption that there are many firms in each nest simplifies the analysis: it implies that firms take the inclusive value as given when taking derivatives of their market share with respect to the wage in their first-order condition. In this case, the elasticity of labor supply to the firm in equation 29 becomes simply

$$\frac{\partial \log s_j}{\partial \log w_j} = \frac{\eta}{\lambda} \tag{31}$$

This is equivalent to setting the market share of the firm to zero within the market, reflecting the assumption that the firm is infinitesimal relative to the size of the labor market. In this sense, this is a model of monopsonistic competition as in section 2.3.1 above, but with discrete choice instead of the representative household framework.

2.3.6 Simultaneous Labor and Product Market Power

In the monopsonistic competition and oligopsony models in the previous sections, we have assumed that the product market was competitive. If we assume instead that firms produce differentiated goods, then firms have market power simultaneously in the labor market *and* the product market. This kind of model is derived in a representative household framework by Deb, Eeckhout, Patel, and Warren (2022), and in a discrete-choice framework by Kroft, Luo, Mogstad, and Setzler (2022).[14]

The simplest model to illustrate how this works is a model of simultaneous monopsonistic and monopolistic competition. We use the same notation as in section 2.3.1 above. Consider a representative household that has preferences over a bundle of differentiated consumption goods C and a bundle of labor supply to differentiated jobs L:

$$U(C, L) = C - \frac{1}{\varphi^{\frac{1}{\varphi}}} \frac{L^{1+\frac{1}{\varphi}}}{1+\frac{1}{\varphi}}, \tag{32}$$

where

$$C = \left[\int_0^1 c_i^{\frac{\sigma-1}{\sigma}} di\right]^{\frac{\sigma}{\sigma-1}} \tag{33}$$

and

$$L = \left[\int_0^1 l_i^{\frac{\theta+1}{\theta}} di\right]^{\frac{\theta}{\theta+1}} \tag{34}$$

Product market power is introduced through the assumption that consumption goods are differentiated, with σ as the parameter that governs the differentiation. When σ goes to infinity, consumption goods become perfect substitutes and the product market becomes perfectly competitive; thus we go back to the case of monopsonistic competition in the labor market without product market power in section 2.3.1. As before in equation 11, when θ goes to infinity, monopsony power disappears. The household's budget constraint is $\int_0^1 p_i c_i di = \int_0^1 w_i l_i di + b$, where p_i is the price of product i, c_i is the consumption of product i, and b is nonlabor income.

Firms take aggregate consumption and aggregate labor as given, and face product demand given by:

$$p_i = \left(\frac{c_i}{C}\right)^{-\frac{1}{\sigma}} P, \tag{35}$$

where $P \equiv \left[\int_0^1 p_i^{1-\sigma} di\right]^{\frac{1}{1-\sigma}}$ is the Dixit-Stiglitz price index. Firm i's inverse labor supply, as before in equation 11, is given by:

$$w_i = \left(\frac{l_i}{L}\right)^{\frac{1}{\theta}} W \tag{36}$$

[14] Azar and Vives (2021) derive a model of simultaneous product and labor market power, in a general equilibrium oligopoly context instead of monopolistic/monopsonistic competition.

Recall that $W = \left[\int_0^1 w_j^{1+\theta} dj\right]^{\frac{1}{1+\theta}}$.

Firm i maximizes profits taking into account the effects of its actions on both its price and its wage. We assume that firm i has a production function that uses only labor, $c_i = F(l_i)$. We express profits in real terms.[15] The firm's profit maximization can be written as:

$$\max_{l_i} \frac{p_i}{P} c_i - \frac{w_i}{P} l_i \tag{37}$$

The first-order condition of firm i sets the marginal revenue from an additional worker equal to marginal revenue product of labor (MRPL, or $R'(l_i)$ in the notation of equation (1)) to the marginal cost of labor (MCL):

$$\underbrace{\left(1 - \frac{1}{\sigma}\right) \frac{p_i}{P} F'(l_i)}_{MRPL} - \underbrace{\left(1 + \frac{1}{\theta}\right) \frac{w_i}{P}}_{MCL} = 0 \tag{38}$$

Alternatively, we could divide this expression by $F'(l_i)$ and it would read "marginal revenue (MR) equal to marginal cost (MC)" (and this is the first-order condition that obtains directly when the problem is rewritten for the firm to choose its output level instead of employment l_i):

$$\underbrace{\left(1 - \frac{1}{\sigma}\right) \frac{p_i}{P}}_{MR} - \underbrace{\left(1 + \frac{1}{\theta}\right) \frac{w_i}{P} \frac{1}{F'(l_i)}}_{MC} = 0 \tag{39}$$

The equilibrium markdown is the wedge between the MRPL and the real wage, and it is given by:

$$\frac{R'(l_i) - \frac{w_i}{P}}{\frac{w_i}{P}} = \frac{1}{\theta} \tag{40}$$

Interestingly, this is the exact same markdown formula as in the case of no product market power (see equation 10). However, note that in this case the MRPL $R'(l_i)$ includes the firm's product market power. Thus, in this case, unlike in equation 10, MRPL is different from the value of the marginal product of labor (VMPL), which is defined as the price that the firm receives from selling its product times the number of additional units that it produces when it hires an extra worker. The formula for VMPL in this model is $\frac{p_i}{P} F'(l_i)$. The formula for the MRPL includes the term $1 - \frac{1}{\sigma}$ to take into account the fact that, when a firm hires an extra worker and increases output, this reduces its price, which has a negative effect on marginal revenue.

The equilibrium markup is the wedge between the relative price of the firm and the marginal cost (expressed relative to the price level), given by:

$$\frac{\frac{p_i}{P} - MC}{\frac{p_i}{P}} = \frac{1}{\sigma} \tag{41}$$

[15] It is helpful to express profits in real terms in order to allow for the general case where firms can influence the overall price level in the economy P. In the case when the firms take the price level as given, as in the model presented here, it does not matter whether profits are expressed in nominal or real terms. However, it does matter in the more general case in which firms can affect the price level in the economy: in that case, if the profits were expressed in nominal terms, the equilibrium would depend on the choice of price normalization (Gabszewicz and Vial, 1972; Azar and Vives, 2021).

Under perfect competition in both product and labor markets, the wage is equal to the value of the marginal product of labor, which is also equal to the marginal revenue product of labor. In this model with labor and product market power, we have two forces that create two wedges between the wage and the value of the marginal product of labor:

1. Product market power creates a wedge (markup) between the *value* of the marginal product of labor and the marginal *revenue* product of labor.

2. Labor market power creates a wedge (markdown) between the marginal revenue product of labor and the wage.

In the monopsonistic competition model of the previous section (and also in the oligopsony model), we only had the second force, driven by labor market power.

The overall effect of product *and* labor market power is summarized in the following formula for the equilibrium gap between VMPL and the wage:

$$\frac{\frac{p_i}{P}F'(l_i) - \frac{w_i}{P}}{\frac{w_i}{P}} = \frac{1+\frac{1}{\theta}}{1-\frac{1}{\sigma}} - 1 = \frac{\frac{1}{\theta}+\frac{1}{\sigma}}{1-\frac{1}{\sigma}} \tag{42}$$

Note that, when σ goes to infinity, product market power goes away, and so does the product market power component of the gap (i.e., the denominator goes to one). Similarly, when θ goes to infinity, labor market power goes away. Conversely, when σ decreases, i.e. product market power increases, workers get paid less relative to the value of their marginal product. In this sense, product market power reduces wages beyond the effects of labor market power.

2.4 Search frictions

In search and matching models, the key deviation from perfect competition is search costs, as well as workers' imperfect information. Firms can pay workers less because workers cannot instantaneously meet all available jobs. One example of how imperfect information affects the labor market is the following: workers do not know which jobs other workers have already applied to, and so several workers may end up randomly competing for the same job. Therefore, even when there are more vacancies than job seekers, there is no guarantee that an individual worker will get hired. This assumed lack of coordination is a key mechanism introducing frictions in search and matching models.

2.4.1 Monopsony power with infinitesimal firms: random search, wage posting, and on the job search

Burdett and Mortensen (1998) build a search model in which firms post wages and have monopsony power despite being infinitesimal in the labor market. If workers search for jobs both when they are employed and when they are unemployed, the equilibrium is characterized by a distribution of wages as opposed to a single wage, even when all firms and workers are identical.

In the Burdett-Mortensen model, all firms are infinitesimal relative to the labor market. There is a continuum of workers with mass M_w, and a continuum of firms with mass M_f.

A firm in the labor market offers a wage w, taking as given the distribution of wages across firms in the labor market $F(w)$. Both employed and unemployed workers receive offers, at an exogenous

Poisson arrival rate of λ. There is also an exogenous rate of separation to unemployment δ. The productivity of an employed worker is A, and the flow of income in unemployment for the worker is b.

In steady state, the number of workers going into unemployment must be the same as the number of workers leaving unemployment. This implies that $\lambda u M_w = \delta(1-u)M_w$. Therefore, the steady state unemployment rate is:

$$u = \frac{\delta}{\delta + \lambda} \tag{43}$$

If a firm increases its wage, then (i) it becomes more successful at hiring the workers it meets, and (ii) it experiences a decrease in the probability that its current employees leave when they receive an offer from other firms. Denote $N(w; F)$ the steady-state level of employment of a firm with wage w when the distribution of wages across firms is F. Then, a firm that sets a wage of w has steady-state profits equal to

$$\pi(w; F) = (A - w)N(w; F) \tag{44}$$

How do we obtain an expression for $N(w; F)$? If a firm pays a wage w, it will lose workers when they go to unemployment, and when they receive offers from other firms that pay a wage higher than w. Thus, if a firm pays w, its separation rate is:

$$s(w; F) = \delta + \lambda(1 - F(w)). \tag{45}$$

The total number of workers the firm loses is equal to the product of the separation rate and the firm's level of employment: $s(w; F)N(w; F)$.

The firm also receives a flow of recruits. If a firm pays a wage w, it recruits a share of the unemployed workers that randomly match with the firm, as well as employed workers that randomly match with the firm, if their current firm pays less than w. Mathematically, the flow of unemployed workers to the firm is $\lambda u \frac{M_w}{M_f}$, and the flow of employed workers to the firm is $\lambda(1-u)G(w; F)\frac{M_w}{M_f}$, where $G(w)$ is the distribution of wages across workers.

The distribution of wages across workers in the labor market is different from the distribution of wages across firms in the labor market. The reason is that higher wage firms have more workers. To obtain an expression for $G(w)$, consider the set of workers with wages less than or equal to w. There can be transitions among workers between firms in this set, which doesn't affect the fraction of workers in the set. However, there cannot be transitions from workers outside of this set directly to this set, because those workers' wages are higher than w and they would not accept offers from firms paying w or less. Thus, the net transitions into the set are only from unemployment. The rate of transitions into this set is the fraction of unemployed workers who get an offer $u\lambda F(w)$. Transitions out of this set can be both to firms with wages higher than w, at a rate $\lambda(1-u)G(w)(1-F(w))$, or to unemployment, at a rate $\delta(1-u)G(w)$.

In steady state, the rate of entry and the rate of exit from this set have to be equal, implying a distribution of wages across workers:

$$G(w; F) = \frac{\delta F(w)}{\delta + \lambda(1 - F(w))} < F(w) \tag{46}$$

With this expression for the distribution of wages across workers, we can also now express the flow of recruits to the firm as:

$$R(w;F) = \frac{M_f \lambda \delta}{M_w} \left(\frac{1}{\delta + \lambda(1 - F(w))} \right) \tag{47}$$

In steady-state, the flow of recruits to the firm has to be equal to the flows of workers out of the firm, that is $R = sN$. This gives us an expression for $N(w; F)$:

$$N(w;F) = \frac{M_f \lambda \delta}{M_w \left\{ \delta + \lambda \left[1 - F(w) \right] \right\}^2} \tag{48}$$

Substituting into the profit function implies:

$$\pi(w;F) = \frac{M_f \lambda \delta (A - w)}{M_w \left\{ \delta + \lambda \left[1 - F(w) \right] \right\}^2} \tag{49}$$

It can be shown that in equilibrium, firms are indifferent between setting any wage between b and $A - \left(\frac{\delta}{\delta+\lambda}\right)^2 (A - b)$. They can choose a lower wage and have higher profit margins per worker, but have a lower level of employment, or they can choose a higher wage and make lower profit margins per worker, but at a higher level of employment. The equilibrium distribution of wages across firms adjusts such that there is indifference between any two wages. The equilibrium distribution of wages across firms is:

$$F^*(w) = \frac{\delta + \lambda}{\lambda} \left(1 - \sqrt{\frac{A - w}{A - b}} \right) \tag{50}$$

Because firms have the same productivity and are indifferent across wage levels, the firm-level markdown is not uniquely defined (this is why there is no entry for the firm-level markdown in our summary Table 1).

The average wage is in between the utility from unemployment and the value of the marginal product of labor:

$$E(w) = \frac{\delta}{\delta + \lambda} b + \frac{\lambda}{\delta + \lambda} A \tag{51}$$

When unemployment is higher, the wage is lower and closer to the utility of unemployment, and when unemployment is lower, the wage is closer to the value of the marginal product of labor. If the unemployment rate goes to zero (which happens when the job separation rate δ goes to zero), the wage goes to the value of marginal product of labor A. If the job offer arrival rate λ goes to infinity, then the wage also goes to the marginal product of labor. This shows that, in this model, monopsony power is enabled by both a non-zero job loss probability, and a less than one probability of job finding.

One can write the markdown relative to the expected wage $(A - E(w))/E(w)$ as follows:[16]

$$\frac{A - b}{b + \frac{\lambda}{\delta} A} \tag{52}$$

[16]Note that this is not exactly the same as the expected markdown.

When firm productivity is heterogeneous and uniformly distributed between zero and one, Burdett and Mortensen (1998) derive an expression for the equilibrium wage as a function of productivity:

$$w(A) = \frac{\frac{\lambda}{\delta}A^2}{1+\frac{\lambda}{\delta}} \tag{53}$$

This implies that the firm-level markdown is:

$$\frac{1+\frac{\lambda}{\delta}(1-A)}{\frac{\lambda}{\delta}A} \tag{54}$$

Postel-Vinay and Robin (2002) consider a version of Burdett and Mortensen (1998) where firms are able to counter-offer when employed workers receive an outside offer. Indeed, in Burdett and Mortensen (1998), firms are passive, and workers leave when they receive a higher offer than their current wage. In Postel-Vinay and Robin (2002), when an employed worker receives an offer from another firm, the firms engage in Bertrand competition. When all workers have the same opportunity cost of employment and all firms are also identical, there are only two wages in this economy: the "monopsony wage" offered to the unemployed, and the competitive wage offered to the employed who received an outside job offer and who benefited from the all out Bertrand competition between the two firms. When introducing firm productivity dispersion, there is additional wage dispersion, because more productive firms are able to offer higher wages to poach already employed workers. Offers and counter-offers for employed workers thus play an important role in wage growth, and allow workers to escape the "monopsony wage."

2.4.2 Monopsony power with Large Firms: Labor Market Concentration in a Random Search Model

Jarosch, Nimczik, and Sorkin (2019) develop a random search model with multi-vacancy firms, based on the Diamond-Mortensen-Pissarides wage bargaining framework. In contrast to the approach in Burdett and Mortensen (1998), where firms post wages, in the Jarosch, Nimczik, and Sorkin (2019) framework, firms do not post wages and bargain over wages with workers after meeting them. The key difference between Jarosch, Nimczik, and Sorkin (2019) and the standard random search and bargaining model is that, in Jarosch, Nimczik, and Sorkin (2019), firms are not atomistic, so there is a positive probability that a worker will encounter the same firm in the future.

Consider a model with N firms in which the firms are granular, i.e., large relative to the market. Job openings occur exogenously, and the probability that a job opening is by firm i is f_i, where $\sum_{i=1}^{N} f_i = 1$. We refer to f_i as the firm's market share or its "size". There is a fixed cost c_i per job opening. The productivity of a worker is normalized to one.

The matching process involves u workers applying for v vacancies, and the job finding rate is $\lambda \equiv \frac{v}{u}(1-e^{\frac{v}{u}})$. Separations occur exogenously at a rate δ. Firms and workers engage in Nash bargaining, with worker bargaining power parameter $\alpha \in [0,1]$. The value of unemployment is U, and the value of employment at firm i is W_i. The value of unemployment is given by:

$$U = b + \beta\left(\lambda\sum_{i=1}^{N} f_i W_i + (1-\lambda)U\right), \tag{55}$$

where b is the flow of income to an unemployed worker, and β is the workers' discount factor.

The value of working for firm i is given by:

$$W_i = w_i + \beta\left(\delta U + (1-\delta)W_i\right) \tag{56}$$

The main departure from the standard model is that, if bargaining between a firm and an unemployed worker breaks down, the firm will not hire the worker if the worker matches with the firm again during the same unemployment spell (unless the worker is the only applicant, which happens with probability $\underline{\lambda} \equiv e^{-\frac{u}{v}}$). After the worker finds a new job, the no-rehire policy ends: that is, if the worker finds at least one new job and then becomes unemployed again, the firm does not refuse to hire the worker even if the worker rejected the firm's offer in a prior unemployment spell. In an off-equilibrium path in which the bargaining between an unemployed worker and firm i breaks down, the continuation value for the worker is:

$$U_i = b + \beta\left(\lambda\sum_{j\neq i} f_j W_j + \underline{\lambda} f_i W_i + (1-\lambda(1-f_i) - \underline{\lambda} f_i)U_i\right) \tag{57}$$

For this punishment mechanism to work, the firm has to be able to track applicants, and the threat not to hire re-applicants has to be credible. Jarosch, Nimczik, and Sorkin (2019) show that, even if firms cannot commit, the firm's threat not to rehire the worker is still operative as long as it is more costly to the worker than to the firm, which occurs when the firm can find a close enough substitute for the worker.

The value of a bilateral relationship for firm i is:

$$J_i = 1 - w_i + \beta(1-\delta)J_i \tag{58}$$

The value of a job opening is:

$$V_i = -c_i + \beta(1 - e^{\frac{u}{v}})J_i \tag{59}$$

The joint surplus from a match is:

$$S_i \equiv W_i - U_i + J_i \tag{60}$$

The surplus is split between the worker and the firm, such that the worker's surplus is $W_i - U_i = \alpha S_i$, and the firm's surplus is $J_i = (1-\alpha)S_i$.

Jarosch, Nimczik, and Sorkin (2019) show that, in this model, the average wage $\overline{\omega}$ is negatively related to labor market concentration C, in the following way:

$$\overline{\omega} = 1 - (1-\alpha)\frac{1-\beta(1-\delta)}{1-\beta\left(1-\lambda\alpha(1-C)-\delta(1-\tau C)\right)}, \tag{61}$$

where

$$\tau = \alpha\frac{\beta(\lambda-\underline{\lambda})}{1-\beta(1-\lambda)} \tag{62}$$

$$C \equiv \frac{\sum_{k=2}^{\infty}\tau^{k-2}f^k}{1+\tau\sum_{k=2}^{\infty}\tau^{k-2}f^k}, \tag{63}$$

where, in turn, $f^k \equiv \sum_i f_i^k$, such that f^1 is equal to 1, and f^2 yields the labor market HHI. Higher orders of f^k are non-standard measures of concentration that differ from the HHI.[17] In practice, concentration C is very similar to the standard HHI in the Austrian data used by Jarosch, Nimczik, and Sorkin (2019). Here, the HHI measures how often a worker who searches randomly encounters a job vacancy from their current employer: intuitively, the higher an employer's market share, the more likely the worker would meet the employer again and thus not be rehired.

Firm-level wages for two firms i and j depend on their sizes f and are characterized by:

$$\frac{1-w_i}{1-w_j} = \frac{1-\tau f_j}{1-\tau f_i} \tag{64}$$

This model demonstrates that HHI can measure market power not only in an oligopsony model, but also in a search and matching model. In both models, firms with higher market shares have more market power. In the oligopsony model, larger firms face higher wage costs when hiring an additional worker, because they have to raise the wages of their existing inframarginal workers. Therefore, larger firms keep wages low. In Jarosch, Nimczik, and Sorkin (2019), larger firms can keep wages low because they can prevent job seekers from accessing their future vacancies. Therefore, larger firms have more leverage in bargaining with workers over the wage. To note, the distribution of firm size is taken as given in Jarosch, Nimczik, and Sorkin (2019), and therefore the model is not helpful to understand the impact of monopsony power on employment or the firm size distribution. In contrast, it is a useful model to understand workers' bargaining leverage in wage negotiations with firms of different sizes, and to understand how this bargaining leverage may be affected if firms collude.

2.4.3 Monopsony Power in a Directed Search Model where Firms have Multiple Employees

Rudanko (2023) explores wage setting in a directed search model of multiworker firms facing within-firm equalizing constraints on wages. The paper builds a model of directed search as in Moen (1997), but with multi-vacancy firms. When firms have multiple vacancies and equality constraints require them to pay the same wage to their existing workers as they offer new hires, wages are reduced through a mechanism similar to that in the oligopsony model in section 2.2. However, in contrast to the oligopsony model, paying lower wages leads equality-constrained firms to hire *more* workers rather than fewer workers, and equlibrium employment is higher than without equality constraints. The surprising positive employment effect of equality constraints can be explained as follows: lower equilibrium wages encourage firms to post more vacancies, and workers are willing to apply as long as posted wages are higher than the reservation wage.

In this model, there are a large number I of firms. Firm i begins the period with n_i workers, and needs to decide how many vacancies to create v_i, and at what wage w_i. The labor force is normalized to 1. The total number of employed workers is a share $N = \sum_{i=1}^{I} n_i$ of the labor force, and there are $1 - N$ unemployed workers. Output per worker at a firm is constant at A, and the flow income for the unemployed is $b < A$. The cost of posting v vacancies (where v can be non-integer) for a firm with employment n is $\kappa(v/n)n$, where κ is increasing in v/n and convex.

[17] This formula is derived in the Appendix of Jarosch, Nimczik, and Sorkin (2019).

Vacancies that have the same wage are grouped into a labor submarket (possibly composed of just one firm). There is a homothetic matching function $m(v, u)$ that determines how many matches are created as a function of the number of vacancies and unemployed workers (there is no on-the-job search in this model). Tightness in a labor submarket is defined as the number of vacancies per unemployed worker searching in that submarket. The probability that a vacancy is filled in a labor submarket i with tightness $\theta_i = v_i/u_i$ is $q(\theta_i) = m(v_i, u_i)/v_i = m(1, 1/\theta_i)$, while the probability that a worker finds a job in a market with tightness θ_i is $\mu(\theta_i) = m(v_i, u_i)/u_i = m(\theta_i, 1)$. The value of search at a market with tightness θ_i and wage w_i is the probability of finding a job multiplied by the market wage w_i, plus the probability of not finding a job multiplied by b. In equilibrium, the value of search has to be the same across all labor markets and equal to U, such that

$$U = \mu(\theta_i)w_i + (1 - \mu(\theta_i))b \tag{65}$$

Equation (65) implicitly defines tightness as a function of the wage, conditional on the value of search U in the labor market. We denote this function $g(w_i; U)$. This function captures the fact that if firm i increases its wage, then more unemployed workers are attracted to the market, and therefore tightness is lower. Lower tightness in turn implies a higher probability of filling vacancies for the firm, and a lower probability of finding a job for the workers searching in the labor submarket. The firm is aware of its effect on tightness when it chooses its wage and how many vacancies to create. In particular, the firm chooses its wage and the number of vacancies to maximize profits:

$$\max_{w_i, v_i} (n_i + q(\theta_i)v_i)(A - w_i) - \kappa(v_i/n_i)n_i, \tag{66}$$

subject to the constraint in equation (65).

This maximization problem implicitly imposes an equality constraint, namely that the wages of existing employees must also adjust to be equal to the wages of the newly posted vacancies.

This problem is scale-independent and can be rewritten in terms of $x = v/n$, which is the number of vacancies divided by the stock of employed workers of a firm:

$$\max_{x_i} (1 + q(\theta_i)x_i)(A - w_i) - \kappa(x_i) \tag{67}$$

The first-order condition with respect to vacancies is:

$$\kappa'(x) = q(\theta)(A - w) \tag{68}$$

On the left-hand side, we have the marginal cost of increasing vacancies per existing employee. On the right-hand side we have the probability of filling the vacancy multiplied by the profit margin per worker.

The first-order condition with respect to the wage is:

$$1 + q(\theta)x = q'(\theta)g_w(w; U)x(A - w) \tag{69}$$

On the left-hand side, we have the cost of increasing wages by one dollar, which is the expected number of employees after hiring (expressed per existing employee). On the right-hand side, we have the benefit of increasing wages, which is the increase in the probability of filling vacancies when tightness goes down $q'(\theta)$, multiplied by the change in tightness when the wage changes

$g_w(w;U)$, multiplied by the number of vacancies per existing employee, multiplied by the profit margin.

If we don't impose the equality constraint that wages for existing employees have to adjust to the newly posted wages, then the first-order condition for vacancies is exactly the same, and the first-order condition for wages is:

$$q(\theta)x = q'(\theta)g_w(w;U)x(A-w), \tag{70}$$

reflecting that the cost of increasing wages is lower compared to equation 69, because new wages do not apply to the existing workers.

An equilibrium in this model is a wage, a level of vacancies, a level of tightness, and a value of search that solve the maximization problem of the firm, and such that each unemployed worker applies to one firm:

$$1 - N = xN/\theta \tag{71}$$

In both the constrained and unconstrained cases, the equilibrium wage can be expressed as a convex combination of the flow of income when unemployed and the output per worker when employed:

$$w = \gamma b + (1-\gamma)A \tag{72}$$

Note that the expression for the wage is similar to that for the expected wage in the random search model of Burdett-Mortensen (equation 51), in that the wage is a weighted average of the flow of income when unemployed and productivity A. Of course, the expression for the weights is different in the two models.

The equilibrium markdown in this model can be written as:

$$\frac{A-w}{w} = \gamma\frac{A-b}{\gamma b + (1-\gamma)A} \tag{73}$$

In the unconstrained case, $\gamma = \epsilon$, where ϵ is the matching function elasticity:

$$\epsilon = \frac{\mu'(\theta)\theta}{\mu(\theta)} \tag{74}$$

In the constrained case, the expression for γ is:

$$\gamma = \frac{\epsilon/\tau}{1-\epsilon+\epsilon/\tau}, \tag{75}$$

where $\tau = \frac{q(\theta)x}{1+q(\theta)x}$ is the firm's new workers as a fraction of its overall employment (after recruiting). If τ goes to one, which means that all the employment of the firm is composed of new recruits, then γ goes to ϵ, which takes us back to the unconstrained case. If τ goes to zero, then γ goes to one, and the wage goes down to the flow of income when unemployed b. Intuitively, when the number of new recruits is small relative to the firm's existing workforce, the cost of increasing wages to attract more recruits is very high.

Note that γ in the constrained case is between ϵ and one, and is thus always higher in the constrained case than in the unconstrained case. This also implies that the equilibrium markdown is higher, and the equilibrium wage is always lower in the constrained case relative to the unconstrained case (also see the wage equation 72).

The first-order condition for the wage is essentially the same as the first-order condition with respect to the wage in the classical monopsony model (equation 1). In Rudanko (2023), when ϵ – the elasticity of the probability of finding a job with respect to tightness – is higher, the markdown is higher. The reason is that an increase in ϵ reduces the absolute value of the elasticity of the probability of filling out a vacancy with respect to tightness, which is $1 - \epsilon$. This elasticity is directly related to the firm-level elasticity of labor supply in this model, and a lower elasticity of labor supply implies a higher markdown in equilibrium, just as in the classical monopsony model (see equation 2).

Because wages are lower in the constrained case than in the unconstrained case, firms optimally post more vacancies in the constrained case. This can be seen using equation (71), which implies that $\kappa'(x)/q(\theta)$ is increasing in x. Knowing this, we can see from equation (68) that an increase in $A - w$ implies an increase in equilibrium vacancy creation x. Given that workers accept jobs as long as the wage is above the flow of income in nonemployment b, an increase in vacancies in this model leads to higher equilibrium employment.

The intuition for a positive equilibrium employment effect is thus similar to that in the generalized oligopsony model of Azar, Qiu, and Sojourner (2021): the firm has two decisions, a wage decision and a hiring intensity decision. The lower equilibrium wage from monopsony raises the incentive to increase employment, by increasing hiring effort as in Azar, Qiu, and Sojourner (2021) or by increasing the number of vacancies as in the Rudanko (2023) model. It is interesting, however, that in the case of directed search à la Rudanko (2023), there is *always* a positive effect of monopsony power on employment, while in the generalized oligopsony model, the effect is ambiguous. The reason for the unambiguous result in Rudanko (2023) is that firms do not anticipate that increasing vacancies will increase market tightness, because firms see themselves as infinitesimal relative to the labor market. In both cases, however, it is crucial for the positive employment effect for firms to be able to hire from a pool of unemployed workers, as opposed to the classical monopsony case in which the equilibrium is always on the aggregate labor supply curve, and therefore there is no unemployment.

There are two opposing forces at play in these models. The first is the upward-sloping market-level labor supply, which implies that when wages are lower, workers supply less labor to the market. The second force is firms' incentives to invest in costly hiring efforts: when firms have a hiring intensity (Davis, Faberman, and Haltiwanger, 2012) or a vacancy margin, lower wages imply a higher incentive to invest in hiring. This second force can sometimes dominate the first one, leading to an increase in employment when wages are lower. In the case of Rudanko (2023), the first force is shut down because the market-level labor supply is vertical in the relevant wage range (workers accept any wage above the reservation wage). Therefore, in that model with a wage equality constraint and monopsony power, the second force always dominates and employment is higher.

A positive employment effect of monopsony power is an interesting theoretical finding because (i) it helps rationalize some empirical results that find greater concentration leading to lower wages without lower employment, and (ii) it has very different implications for the effects of monopsony power on efficiency compared to the classical model. In the oligopsony model, equilibrium employment is below the efficient level, while in the Rudanko (2023) model, equilibrium employment is above the efficient level. How could employment be inefficiently high? An excessive employment equilibrium happens when workers have high employment but low wages: in this case, average workers' expected income is lower than in the constrained efficient equilibrium that has lower employment but significantly higher wages. In other terms, it can be more advantageous for workers to cycle more often through unemployment if this maintains sufficiently high wages in equilibrium.

2.5 Discussion

2.5.1 The Role of Firm Size Heterogeneity

Our theoretical discussion makes it clear that monopsony power can exist even when firms all have very small market shares, due to e.g. job differentiation and search frictions.

While the core theoretical insights about the sources of monopsony power do not hinge on the explicit consideration of firm size heterogeneity within a market, incorporating this aspect reveals additional, nuanced insights into labor market outcomes, monopsony power, and firm size. Here, we use firm size and firm share interchangeably, but strictly speaking the theoretical results are about firm share within a labor market.

Firstly, differences in firm size within a labor market affect the formula for the HHI. When some firms are larger than others, the HHI is no longer just the inverse number of firms, but reflects the firm size distribution.

Secondly, differences in firm size also lead to markdown heterogeneity within a market; whether larger firms have lower or higher markdowns depends on the model. In oligopsony models and job differentiation models, the markdown increases in firm share (see Table 1). By contrast, in Burdett and Mortensen (1998), larger firms pay higher wages and make lower profits per worker.

Thirdly, when firm size differences are endogenized based on productivity, complex relationships emerge between the markdown, the wage, productivity, and firm size. In Burdett and Mortensen (1998), the markdown still decreases with firm productivity and size, at least as long as productivity is uniformly distributed (see equation 54 above). However, in Postel-Vinay and Robin (2002); Rudanko (2023), the relationship is more complex. The rise in markdown with firm share and productivity dampens the rise in wage with productivity, and, under some conditions, the relationship could be reversed, resulting in more productive firms paying lower wages. For example, in Postel-Vinay and Robin (2002), this reversal can occur because more productive firms can attract new workers with the promise of larger future wage increases, allowing more productive firms to offer lower wages to new recruits. Therefore, one cannot simply conclude that oligospony models fit the data better than search models from the fact that there is a positive relationship between markdown and firm size.

Lastly, differences in firm productivity and firm size within a labor market can also help explain labor reallocation after a minimum wage increase. We will explore this further in section 4.3, drawing on Dustmann, Lindner, Schönberg, Umkehrer, and vom Berge (2022).

2.5.2 Determinants of monopsony power across models

Class of monopsony model	Average Markdown across Firms	Markdown at firm level	Notes
Oligopsony	$\frac{HHI}{\eta}$	$\frac{s_j}{\eta}$	η is the labor supply elasticity to the market. The employment share of firm j is s_j, and the HHI is based on employment shares.
Differentiated jobs			
Monopsonistic competition	$\frac{1}{\theta}$	$\frac{1}{\theta}$	θ goes to infinity means jobs are perfect substitutes
Oligopsony	$\frac{1}{\theta} + \left(\frac{1}{\varphi} - \frac{1}{\theta}\right) HHI$	$\frac{1}{\theta} + \tilde{s_j} \left(\frac{1}{\varphi} - \frac{1}{\theta}\right)$	φ is the aggregate elasticity of labor supply. The payroll share of firm j is $\tilde{s_j}$, and the HHI is based on payroll shares.
Discrete choice (logit)	$\approx \frac{1}{\eta}\left[\frac{J(J^2-3J+1)}{(J-1)^3} + \frac{J^3}{(J-1)^3} HHI\right]$	$\frac{1}{\eta(1-s_j)}$	η is the coefficient on log wage in the utility, thus related to wage elasticity of labor supply. The number of firms in the labor market J also affects the firm-level wage elasticity of labor supply.
Search and matching			
Burdett-Mortensen	$\frac{A-b}{b+\frac{\lambda}{\delta}A}$		A and b are the productivity of an employed worker, and the income of a non-employed worker, respectively. λ/δ is the ratio of the job-finding rate and the separation rate.
Rudanko	$\gamma\frac{A-b}{\gamma b+(1-\gamma)A}$ (if firms homogeneous)	$\gamma\frac{A-b}{\gamma b+(1-\gamma)A}$	γ is a function of the matching function elasticity, and takes different values in the constrained and unconstrained case.

Table 1: Markdowns across models (see section 2 for more details)

Class of monopsony model	Firm-level elasticity	Market-level elasticity	Notes
Oligopsony	$\frac{\eta}{s_j}$	η	η is the labor supply elasticity to the market. The employment share of firm j is s_j.
Differentiated jobs			
Monopsonistic competition	θ	φ	θ goes to infinity means jobs are perfect substitutes. Firms do not take into account φ when setting wages.
Oligopsony	$\frac{\theta\varphi}{\varphi+\tilde{s_j}(\theta-\varphi)}$	φ	θ goes to infinity means jobs a perfect substitutes. The payroll share of firm j is $\tilde{s_j}$.
Discrete choice (logit)	$\eta(1-s_j)$	0	η is the coefficient on log wage in the utility.

Table 2: Elasticity of labor supply across models (see section 2 for more details)

Theoretically, the wage markdown increases (and hence wages decrease) with (see Table 1):

- The inverse labor supply elasticity
- Labor market concentration
- The inverse substitutability between jobs
- The opposite of the income flow while unemployed

Recent empirical work on monopsony power has focused on estimating the labor supply elasticity and the impact of labor market concentration on wages. Both of these quantities can affect wages across a variety of models. Labor market concentration becomes particularly relevant when there are relatively few firms in the labor market, as would be the case in markets affected by anticompetitive mergers.

Table 3 organizes the various types of models and mechanisms of monopsony power documented in the literature. For each class of models, we summarize the key deviation from the perfectly competitive labor market model, and in particular what gives firms the ability to pay wages below the marginal revenue product of labor (marginal productivity). In the last column, we summarize the key factors that determine the level of labor market power in each class of models. The Table can help guide researchers in choosing modeling frameworks that best capture key features of the problem they are interested in. Oligopsony models are the most straightforward to derive, and are particularly useful when analyzing changes in firm shares in the labor market, for example as a result of mergers. Search models are useful to analyze job search, job finding and recruitment and retention efforts, e.g. when evaluating the effects of non-competition agreements. Finally, differentiated job models are particularly useful when dealing with non-wage job characteristics and amenities, such as geographic distance. Our policy-oriented monopsony review paper (Azar and Marinescu, forthcoming) illustrates how these models may be applied to a large array of policy questions.

3 Empirically measuring monopsony power

3.1 Definition of the markdown in the empirical literature

Based on the theories in section 2, defining the markdown as $(MRPL - w)/w = MRPL/w - 1$ is most practical. This definition directly connects to the formulas in Table 1, facilitating a smoother transition from the empirical estimate of the markdown to the theoretical framework. The markdown $(MRPL - w)/w$ measures the percent wage increase that would result, all other things equal, if all monopsony power were eliminated so that $w = MRPL$. We will use this definition of the markdown in our tables that summarize empirical estimates (Tables 4, 5 and 6). It is important to know that the empirical literature has also used other definitions of the markdown, including $MRPL/w$, or $(MRPL-w)/MRPL = 1 - w/MRPL$, or $w/MRPL$. All of these alternative definitions of the markdown, except the last one ($w/MRPL$), are monotonically increasing functions of the markdown as we define it. The reader should be especially careful when an article defines the markdown as $w/MRPL$ because this expression is *decreasing* in the other definitions of the markdown, and this means that the signs of relationships with other variables

Class of monopsony model	Key deviation from perfect competition	Why firms can pay workers less than their marginal productivity	Determinants of market power
Oligopsony	Finite number of firms that engage in strategic interactions	Workers' labor supply elasticity to the market is limited	Labor market concentration, labor supply elasticity to the market
Search	Matching frictions	Workers cannot instantaneously meet all available jobs	Labor market concentration, workers' outside option, matching frictions, bargaining power (random search)
Differentiated jobs with monopsonistic competition	Heterogeneous preferences over jobs that differ in wages and amenities	The job is better than the next best alternative for the worker	Job substitutability (workers' preferences)

Table 3: Mechanisms of monopsony power across models

are opposite. In general, readers should double-check what definition of the markdown authors are using.

3.2 Elasticity of labor supply

Authors typically calculate the markdown as the inverse of their estimated firm-level labor supply elasticity, which is consistent with most models, as can be seen in Tables 1 and 2.

Identifying the elasticity of labor supply requires a plausibly exogenous source of variation in the wage. Further, it is important to consider whether one is measuring the elasticity of labor supply to an individual firm or to a whole labor market (for theoretical implications of the market-level vs. firm-level elasticities see Tables 1 and 2). If the elasticity of labor supply is to the market, then one needs to define which market this is.

Further, the empirical setups used to identify the labor supply elasticity may focus on a specific part of the process leading to a firm's employment level. Some papers focus on job applications and/or hires, while others focus on job separation. This focus may be due to data availability, and/or may enable a stronger identification strategy. Whatever the reasons for this restricted focus, when a paper does not have as an outcome the firm's employment level, one needs to transform the elasticity estimates in order to get to a labor supply elasticity (wage effect on employment levels) rather than e.g. an application elasticity or a separation elasticity. The authors typically perform these transformations, and we report below the labor supply elasticity as calculated by the authors.

3.2.1 Experiments

Dal Bó, Finan, and Rossi (2013) examine an experiment conducted as part of Mexico's Regional Development Program (RDP) to increase the federal presence in certain municipalities. In the first stage, two different wage offers were randomly offered to prospective public servants across recruitment sites. In the second stage, eligible applicants were selected at random to receive job offers. A 33% increase in wages led to a 26% increase in applications and a 35% increase in the conversion rate (i.e. the rate a which job offers are converted into filled vacancies), implying a labor supply (arc-)elasticity of around 2.15.

Dube, Jacobs, Naidu, and Suri (2020) use Amazon Mechanical Turk (MTurk) data from 2014 to 2017, including data from five MTurk experiments. They use both the experiments, and observational data with machine learning to estimate the labor supply elasticity. They find low labor supply elasticities, with an elasticity of 0.096 based on the double machine learning specification, which is their preferred specification. Experimental estimates give similar elasticities. Based on the range of elasticities estimated, they conclude that workers on MTurk are paid at most 13% of their productivity.

Caldwell and Oehlsen (2022) analyze Uber field experiments with an "earnings accelerator" feature in the U.S., in Houston (2017) and Boston (2016). This accelerator introduces temporary, exogenous variation in wages among Uber ridesharing drivers. The experiments offered drivers a 10%-50% increase in earnings per trip for one week. For a 10% increase in wages, female drivers work 8% more hours (elasticity of .8), and male drivers work only 4% more hours (elasticity of .4). Uber's firm level labor supply elasticity is interpreted as a market level labor supply elasticity when Uber is a monopsonist for drivers who could not shift to Lyft, while Uber's elasticity is interpreted as a firm-level elasticity when drivers could shift to Lyft. Based on this approach, the authors find a market level labor supply elasticity of 0.8 and a firm-level elasticity of 1.5. In the Uber data, there is variation in both hours (intensive margin) and employment (extensive margin), while most of the literature estimates a firm-level labor supply elasticity by focusing on the extensive margin. To allow for a more straightforward comparison, the authors calculate an extensive margin firm-level labor supply elasticity, while accounting for additional driver entry that would occur if Uber were to increase its wages for all drivers. Calculated in this way, the firm-level labor supply elasticity is about 2. Based on the latter estimate, the authors report a markdown of 30%.

3.2.2 Other forms of identification

Bassier, Dube, and Naidu (2022) use Oregon, USA, UI microdata from 2000-2017. They estimate the separation elasticity, i.e. the change in the probability that a worker separates from a firm in response to a change in the wage. They use fixed effects and matching for identification. They match workers with relevant controls by focusing on pairs of workers who began at the same origin firm (step 1) and transitioned to a new firm (step 2). The separation elasticity at the new step 2 firm is estimated by comparing the separation probability of workers who are in a step 2 firm with a higher wage vs. a step 2 firm with a lower wage. They estimate the firm's wage using firm fixed effect derived from an AKM regression (Abowd, Kramarz, and Margolis, 1999), which allows them to focus on the part of the wage that is due to the firm rather than to worker characteristics. They use a stacked event-study design centered on the transition to a new firm (step 2), and where the outcome of interest is workers' separation probability to a final firm or to unemployment. They regress the separation outcome on the individual wage difference between the new (step 2) and old (step 1) firm instrumented by the difference in the AKM firm fixed effects, and controlling for a rich set of variables capturing workers' history (including a fixed effect for the origin firm). The separation elasticity estimate from the preferred matched event study approach is -2.1. The preferred labor supply elasticity is 4.2. They find markdowns of around 20%.

Goolsbee and Syverson (2023) estimate the labor supply elasticity for faculty at US four-year non-profit colleges and universities. They estimate the inverse labor supply elasticity to the firm by regressing faculty wages on faculty employment. They instrument employment with a labor demand shifter, here the lagged number of student applications at the institution. They find a 5.1 firm-level labor supply elasticity for tenure-track faculty and a highly elastic labor supply for non-tenure track faculty (28.6). Among tenure-track faculty, the labor supply elasticity declines with rank, with a labor supply elasticity of 7.7 for assistant professors, 3.0 for associate professors and 1.8 for full professors. This is evidence of monopsony power for tenure-track faculty, while the market for non-tenure track faculty is fairly competitive.

Amodio and de Roux (2023) add to the literature by estimating markdowns for a developing country, Colombia, focusing on the estimation of the firm-level labor supply elasticity for manufacturing firms. They estimate the impact of plausibly exogenous exchange rate shocks separately on average wages and on employment at the plant level. The ratio of these coefficients gives the elasticity of the inverse labor supply curve. Using this strategy, they find a firm-level labor supply elasticity of around 2.5, which corresponds to a markdown of about 40%. Further, they find that the firm-level inverse labor supply elasticity is larger in firms with above median market share, consistent with the predictions of the oligopsony model (recall that, per equation 2 above, the inverse firm-level labor supply elasticity is s_j/η, where s_j is the firm's market share and η is the labor supply elasticity to the market).

The elasticity of labor supply and markdown estimates from the articles discussed in this section are summarized in Table 4.

Table 4: Elasticity of labor supply and markdowns across studies estimating the elasticity of labor supply

Study	**Data set**	**Identification**	**Firm-level elasticity**	**Market-level elasticity**	**Estimated markdown**[a]
Dal Bó, Finan, and Rossi (2013)	Regional Development Program (RDP) experiment in Mexico (2011)	Random assignment of wages and job offers to prospective public servants	2.15		not reported
Dube, Jacobs, Naidu, and Suri (2020)	Observational and experimental Amazon Mechanical Turk (MTurk) data, 2014-2017	Plausibly exogenous variation in pay using double machine learning	0.096		Workers are paid less than 13% of their productivity.[b]
Caldwell and Oehlsen (2022)	Uber data from the US, 2016 (Boston) and 2017 (Houston)	Random assignment of Uber driver pay increases; comparison of drivers with/without access to Lyft to get at market vs. firm level elasticity	2	0.8[c]	30%
Bassier, Dube, and Naidu (2022)	Oregon (USA) unemployment insurance data, 2000-2017	Compare the worker separation probability between high and low wage firms using fixed effects and matching	4.2		20%
Goolsbee and Syverson (2023)	IPEDS data on faculty at US four-year not-for-profit colleges and universities, 2002-03 to 2016-17	Estimate inverse labor supply elasticity using lagged college applications as instrument for faculty employment	5.1[d], 28.6[e]		not reported
Amodio and de Roux (2023)	Colombian manufacturing, exports, and exchange rate data, 1994-2009	Response of wages and employment to plant-specific shocks due to plausibly exogenous exchange rate variation	2.5		40%

[a] All estimated markdowns are reported as percentages calculated as (MRPL-w)/w, where MRPL is marginal revenue product of labor and w is the wage.
[b] The markdown corresponding to an elasticity of 0.096 would be 1/0.096=1042%.
[c] Market defined as ride-share driving in Houston; this elasticity includes both extensive and intensive margins.
[d] For tenure-track faculty.
[e] For non-tenure track faculty.

3.3 Labor Market Concentration

3.3.1 Concentration, market power and welfare

In this section, we elucidate the intricate relationships between concentration, wages, and overall welfare. Higher concentration is associated with greater market power, but it can also be associated with higher wages if it results from higher productivity, leading to ambiguous welfare effects. Additionally, the effects of concentration on welfare depend on the definition of the social welfare function, and whether one takes into account income inequality and positional externalities, or the potential impacts of market power on the political system.

In oligopsony models (with or without differentiated jobs), there is a positive relationship in equilibrium between the average markdown in a labor market and its concentration level (HHI). Since markdowns are typically not directly observable, empirical work often uses wages as the outcome variable. The variation in wages drives the variation in the markdown ($MRPL/w - 1$), but only if we can hold productivity $MRPL$ fixed. Further, when attempting to explain variation in wages, it is important to remember that factors other than concentration theoretically affect the equilibrium markdown. These factors include the market-level labor supply elasticity and the degree of substitutability between jobs (see Table 1). Therefore, even if, *ceteris paribus*, we expect a negative relationship in equilibrium between wages and concentration, this relationship may not hold empirically when we cannot sufficiently control for the relevant variables.

Miller, Berry, Scott Morton, Baker, Bresnahan, Gaynor, Gilbert, Hay, Jin, Kobayashi, Lafontaine, Levinsohn, Marx, Mayo, Nevo, Pakes, Rose, Rubinfeld, Salop, Schwartz, Seim, Shapiro, Shelanski, Sibley, Sweeting, and Wosinska (2022) illustrate the concerns that industrial organization economists have with regressing prices or wages[18] on concentration.[19] They demonstrate through a theoretical model that higher concentration can coexist with lower prices when there are also productivity differences across markets. This insight is pivotal in understanding that the relationship between market concentration and price levels is not always positive, and therefore one may not find empirically that concentration increases prices. Consider two firms competing in a Cournot duopoly in market A, both with equal costs and market shares. These firms also compete in market C, where one firm enjoys lower costs due to higher productivity, thus gaining a larger market share. The result is a paradoxical scenario where market C, despite higher concentration, has lower prices than market A. This example underscores that higher productivity, leading to lower average costs, can result in lower prices even if it makes the market more concentrated. However, these lower prices do not imply that *market power* is reduced, nor that concentration is a poor measure of market power. Rather, the example highlights that market power, defined as the ability to set prices above marginal cost, can increase alongside consumer welfare if firms are more productive.[20]

[18] Remember that wages are the price of labor services.

[19] Miller et al. (2022) note that variation in concentration that arises due to mergers (or entry or exit events) can be informative about the effects of mergers on competition. Thus, to the extent that other confounding factors can be accounted for, a merger is a quasi-experiment that allows for evaluating the impact of HHI on the markdown.

[20] This example also highlights the potential for confusion arising from the label "pro-competitive" in antitrust jargon. The confusion arises from the definition of increased competition: is it lower prices or lower markups? In antitrust parlance, "pro-competitive" refers to a merger that leads to lower prices. Yet, if we assess competition based on markups rather than prices, then there are some "pro-competitive" mergers that are in fact anti-competitive, because they lower prices but increase the markup. This does not mean that markup-reducing pro-competitive

Not only can higher concentration be sometimes associated with lower prices, it can also be associated with higher wages and higher worker welfare. Recasting Miller et al. (2022)'s example within the labor market context reveals that increased labor market concentration can result in higher wages if it is the result of an increase in a firm's labor productivity. Further, *product* market concentration can also increase worker welfare in some situations. For example, Boar and Midrigan (2024) develop a model where the welfare-maximizing product market regulation simultaneously increases product market concentration, markups, *and* wages. This counterintuitive outcome arises because such regulation shifts production to the most productive firms, which are also the largest firms.

While market power does not invariably diminish welfare, it has the potential to do so, particularly when stemming from anticompetitive mergers (see section 4.1). Additionally, market power can affect the political system, potentially undermining welfare through various channels (Cowgill, Prat, and Valletti, 2023; Khan and Vaheesan, 2017; Wu, 2018; Zingales, 2017). Further, the utilitarian welfare function may not fully capture the broader impact of market power. For instance, even if wages increase, a reduction in the labor share due to heightened market power can erode workers' relative societal status, with negative welfare implications if positional externalities matter (Frank, 2005). Finally, understanding the nuances of labor market power is crucial for evaluating the effectiveness of economic policies, including minimum wage legislation (see section 4.3) and unionization efforts, among others (Azar and Marinescu, forthcoming).

In summary, market power is important for reasons that go beyond its impact on prices and wages. To the extent that labor market concentration captures market power in the labor market, it can be a useful addition to the labor economist's toolbox.

3.3.2 Measurement of labor market concentration

Papers that focus on the impact of HHI on labor market outcomes must not only find plausibly exogenous variation in HHI, but also confront some additional choices:

- Define the labor market: this is necessary in order to calculate the HHI, since the HHI depends on *shares*. For example, some papers define the labor market using industries, and others using occupations.

- Decide on whether one uses data on employment levels (stocks), or on job vacancy or new hires to represent flows of newly available jobs. This choice mostly affects the level of the HHI, with vacancy flow-based HHIs being typically higher than employment stock-based HHIs because firms do not all have open vacancies during a given time period (Marinescu, Ouss, and Pape, 2021). The flow measurement better reflects market conditions that job seeking workers face at a point in time, but flow and stock measures are highly correlated across markets (Marinescu, Ouss, and Pape, 2021).

For example, Azar, Marinescu, Steinbaum, and Taska (2020) calculate concentration measures for the US based on online job vacancies from the website BurningGlass (now Lightcast). An

mergers could not exist. For example, when two smaller firms merge, they may lower marginal costs and capture market share from a larger competitor, decreasing the equilibrium HHI and average markups. The term "pro-competitive" is arguably more aptly applied to such markup-reducing mergers. See section 4.1 below for a more extensive discussion of mergers and their impacts on the labor market.

updated version of these concentration measures based on the Lightcast data is available for researchers to download (Choi and Marinescu, 2024). The data covers 2007-2021 (except 2008-2009) and is at the quarter by commuting zone by 6-digit SOC code level.

Because there are many ways of defining markets, labor market concentration may be a noisy measure of market power. For example, if the labor market is assumed to be the whole state while workers mostly look for jobs in their commuting zone (Marinescu and Rathelot, 2018), then firms' market shares will be systematically underestimated. The empirical literature and antitrust practitioners have developed tools to deal with this challenge. We will return to market definition when discussing merger policy in section 4.1.

Another reason why the HHI may be a noisy measure of labor market power is that, in practice, firms can be linked through a network of interlocking shareholdings (including both cross-ownership and common ownership).[21] The literature has shown that common ownership in particular has increased in recent decades, in part due to the rise of index funds (Matvos and Ostrovsky, 2008; Harford, Jenter, and Li, 2011; Azar, 2012; Posner, Scott Morgan, and Weyl, 2016; Posner and Weyl, 2018; Backus, Conlon, and Sinkinson, 2021; Azar and Vives, 2021).

The first step in quantifying the effect of common ownership for competition is to calculate the weight that each firm puts on the profits of other firms in its objective function.[22] Given these objective function weights λ_{jk}, we can construct a modified HHI (MHHI), first proposed by Bresnahan and Salop (1986); O'brien and Salop (1999), that quantifies the level concentration including both the market shares and the information from the interlocking shareholding network. In particular, in the case of common ownership, given the labor market shares s_j, the MHHI is:

$$MHHI = \sum_{j=1}^{J} s_j^2 + \sum_{j=1}^{J} \sum_{k \neq j} s_j s_k \lambda_{jk} \tag{77}$$

The first term is the HHI in the market, and the second term is the MHHI delta, which quantifies the increase in concentration in the market from common ownership. All of this analysis presupposes that shareholders are powerful enough that managerial objectives do not enter the objective function of the firm. If managers also have some power, then under some assumptions the formula for the MHHI delta is similar, except that there is a mitigation factor τ_j that reduces the shareholders' λ's (Azar and Ribeiro, 2022). Thus, as shareholders become less powerful (for example, if highly dispersed), the MHHI tends towards the HHI, even if there is common ownership.

To sum up, labor market concentration can be used in some contexts as a proxy for market power, especially when researchers focus on explaining variation in labor market level outcomes.

[21] Cross-ownership refers to direct ownership, by a firm that participates in a market, of shares in a competing firm. Common ownership refers to companies having shareholders in common.

[22] We here give the formula for the weight that firm j puts on the profits of firm k in its objective function (relative to its own profits). When there are I shareholders indexed by i and β_{ij} denotes shareholder i's percent ownership stake in firm j and γ_{ij} denotes shareholder i's percent of control in firm j, the weight is:

$$\lambda_{jk} = \frac{\sum_{i=1}^{I} \gamma_{ij} \beta_{ik}}{\sum_{i=1}^{I} \gamma_{ij} \beta_{ij}} \tag{76}$$

For example, if firm j has only one shareholder, and that shareholder has a 30% ownership stake in firm k, then the λ_{jk} weight is 0.3.

3.3.3 Empirical Evidence on Labor Market Concentration

The empirical literature shows that higher labor market concentration is associated with lower wages (Azar et al., 2020; Jarosch, Nimczik, and Sorkin, 2019; Berger, Herkenhoff, and Mongey, 2022a). Focusing on mergers specifically, the literature shows that mergers that greatly increase labor market concentration decrease wages (Prager and Schmitt, 2021; Arnold, 2019). Prager and Schmitt (2021) focus on US hospital mergers and Arnold (2019) considers mergers in all industries in the US. Prager and Schmitt (2021) construct a measure of hospital concentration in each commuting zone (CZ) using data from the American Hospital Association (AHA) on hospital mergers and acquisitions. They use the HHI to capture the market share of each hospital system in a CZ. They find that mergers decrease wage growth for workers with specialized skills when concentration (HHI) increases greatly: for these specialized workers, mergers reduce wage growth by more than 25%. By contrast, they find no effects for mergers that merge hospitals across CZs rather than within CZs, suggesting that effects are not due to firm-wide policies aimed at reducing labor costs, but rather to changes in labor market competition within a CZ. Arnold (2019) finds that mergers that greatly increase labor market concentration reduce earnings by more than 2%. If the effect were driven by increased product market power and output reductions that lead to decreased labor demand, one would expect more negative wage effects for non-tradable industries whose product market is also local. However, Arnold (2019) finds no difference in effects between industries with tradable and non-tradable goods, implying that the adverse wage effect of mergers is not due to changes in product market power. Finally, Arnold (2019) uses the Cournot oligopsony model (see Section 2.2) together with the empirical estimates to calculate an implied market-level labor supply elasticity of 0.87, and an implied markdown of about 5%. We report these estimates in Table 6, alongside estimates from other studies.

Merger effects in the labor market can help identify models that best fit the facts. In oligopsony models, mergers decrease both wages and employment. However, in search and matching models, mergers could decrease wages due to a worsening of outside options (e.g., Jarosch, Nimczik, and Sorkin, 2019), but employment may not decrease (Rudanko, 2023). Prager and Schmitt (2021) finds that hospital mergers do not reduce employment even when they decrease wages, suggesting that a search and bargaining model may be more appropriate in their specific setting. By contrast, Arnold (2019) uses a sample of all US industries and finds a negative employment impact of mergers that greatly increase labor market concentration, which is compatible with the oligopsony model.

Azar, Qiu, and Sojourner (2022) measure the level of concentration from common ownership in US labor markets and find that concentration has increased substantially over time. Using entries of firms into the S&P 500 index as a source of variation in labor market common ownership, they find that common ownership reduces wages, but increases employment. They argue that this cannot be rationalized by a full-employment oligopsony model with an increasing labor supply, but could be explained by a search model of the labor market in which firms increase hiring efforts when wages are lower (Rudanko, 2023; Gottfries and Jarosch, 2023).

3.4 Reduced-form approach based on workers' outside options

Caldwell and Danieli (2023) introduce the Outside Options Index (OOI), a measure of the jobs available to an individual worker based on the characteristics of the job and the worker. This index of outside options is derived using a matching framework, and resembles the measure derived in

Jarosch, Nimczik, and Sorkin (2019) (see section 2.4.2). While the OOI is similar to the HHI, it can be calculated for each individual worker instead of being calculated at the market level. To calculate the OOI, Caldwell and Danieli (2023) use a matched employer-employee dataset from Germany covering the years 1993-2014. They find that willingness to commute or move is an important factor in explaining variation in outside options; this result is consistent with the importance of geography as an element of job differentiation. To identify the impact of outside options on wages and employment, they compare the experiences of individuals involved in the same mass layoff who had different levels of the OOI (as calculated using pre-layoff characteristics). Relative to peers in the same layoff, a worker with one unit higher OOI (slightly more than one standard deviation) has 10% higher earnings (as a share of pre-layoff earnings) one year after the separation, and is roughly 1.5% more likely to be employed.

3.5 Calibration and simulation

Berger, Herkenhoff, and Mongey (2022a) use US Census Longitudinal Business Database microdata covering 1977-2011 to estimate reduced-form labor supply elasticities exploiting changes in state corporate taxes. The reduced-form labor supply elasticity is around two. They use these estimated labor supply elasticities to calibrate their oligopsony model with job differentiation (see section 2.3.2) and firm differences in productivity. Using the calibrated model, they find an aggregate markdown (which they define as $w/MRPL$) of 0.72 (in our definition of the markdown as $(MRPL - w/w)$, the markdown is 39%), corresponding to a labor supply elasticity of 2.6. Comparing steady states at an aggregate Frisch elasticity of labor supply of 0.50, they measure a welfare loss of 7.6% and an output loss of 20.9% relative to a perfectly competitive labor market. They find that roughly 60% of welfare losses are driven by employment misallocation, as more productive firms keep their employment inefficiently low in order to exercise more monopsony power. 30% of welfare losses are due to pure markdowns, and the remainder is due to the interaction of markdowns with misallocation. The estimates from this study are reported in summary Table 6, alongside estimates from other approaches.

Azar and Vives (2021) calibrate a model of general equilibrium oligopoly with simultaneous product and labor market power, and with common ownership. In this model, the equilibrium markdown is a function of the labor market MHHI and the market-level elasticity of labor supply.[23] They find that the increase in common ownership between 1985 and 2017 implies a decline in the labor share that is similar in magnitude to the one observed over that period in the US.[24]

Jarosch, Nimczik, and Sorkin (2019) use Austrian labor market data that covers the universe of private sector employment in Austria from 1997-2015. To get at the mechanism in their model (see section 2.4.2), they measure "re-encounter" rates over a four-year period: taking transitions of workers who left firm a for firm b and then leave firm b as well, what share of these transitions are returns to firm a after leaving firm b? Calibrating their model using these re-encounter rates, they find that eliminating size-based market power increases the labor share by about 1.3 to 2

[23] In this sentence, we use markdown as we have defined it in this chapter. In the paper, they define the "markdown" differently, in physical units, i.e., relative to the marginal physical product of labor instead of the marginal revenue product of labor.

[24] However, the decline in the labor share in the model would be lower if one added corporate governance frictions that reduced the MHHI, as discussed in Section 3.3.2.

percentage points. Since the labor share is around 2/3, this increase amounts to a 2.3 to 3% increase in wages, depending on the year.

3.6 Structural estimation

Azar, Berry, and Marinescu (2022) estimate a nested logit model using job application data from Careerbuilder.com in the US in April-June 2012. They find direct evidence of substantial job differentiation. Geographic distance is an important component in job differentiation: 18% of the variance in the total utility of a job is explained by the geographic distance (ZIP code to ZIP code) between the job and the job seeker. Without the use of instruments for wages, job applications appear very inelastic with respect to wages. Their preferred instrument for the wage is the predicted wage posted by the firm in markets other than the focal market. Using this instrument results in more elastic firm level application supply curves, with a firm level labor supply elasticity of 4.8. This implies a markdown of about 21%. They define the labor market as an SOC-6 occupation by commuting zone, and find that the implied market level labor supply elasticity is 0.5. Considering variation in this elasticity across markets, they conclude that the market level elasticities are low enough to typically make an SOC-6 occupation by commuting zone a plausible definition of a relevant labor market for the purpose of antitrust analysis (also see below section 4.1 on the hypothetical monopsonist test).

Lamadon, Mogstad, and Setzler (2022) use US matched employer-employee panel data together with business tax filings and worker-level filings for 2001–2015. They estimate a differentiated job model, and specifically a nested model with monopsonistic competition (see section 2.3.5). The model incorporates two-sided heterogeneity, and heterogeneous worker preferences for amenities. To recover the firm level labor supply elasticity, they use instruments. They find a firm-level elasticity of 6.5, and a market-level elasticity of 4.6. Overall, their estimates imply that the marginal revenue product of labor is 15% higher than the wage. Regarding job differentiation, they find that average workers would be willing to pay 13% of their earnings to stay in their current jobs. Additionally, they show that more productive firms offer more amenities – consistent with Sockin (2021) –, and that better workers place a higher value on more amenities.

Roussille and Scuderi (2023) use data on full-time, high-wage engineering jobs, with a discrete choice model to calculate a wage markdown. Their data from Hired.com has a number of institutional features that they exploit to obtain rich estimates of both labor supply and demand without requiring an instrument. They focus on workers based in San Francisco, and estimate labor supply elasticities to the firm of 3.6 to 5.7. Under their preferred monopsonistic model of competition, the average predicted markdown is $31,640 (or 19.5% of productivity), which translates to a 25% markdown under our definition of markdown with the wage in the denominator. Further, relevant to the sources of the markdown, they find significant evidence for job differentiation, with heterogeneous worker preferences over amenities.

Kroft et al. (2022) use matched employer-employee panel data for the US construction industry spanning 2001-2015, and data on procurement auctions. They compare first-time procurement contract winners to non-recipients who had never won contracts. They estimate a firm-specific labor supply elasticity of about 4.1, and that wages are marked down 25-29% relative to the marginal revenue product of labor (using our definition of the markdown). Further taking into account imperfect competition in the product market, they find that wages are marked down

44-49% relative to the *value* of the marginal product of labor (for a theoretical definition of this concept, see section 2.3.6). In other terms, product market power and labor market power both significantly contribute to reducing wages relative to perfectly competitive labor *and* product markets.

Table 5 synthesizes the estimates from the articles discussing structural estimation in this section.

Table 5: Elasticity of labor supply and markdowns across studies using structural estimation

Study	Data set	Identification	Firm-level elasticity	Market-level elasticity	Estimated markdown[a]
Azar, Berry, and Marinescu (2022)	CareerBuilder.com application and vacancy data from the US, April-June 2012	Nested logit model. Instrument for the posted wage using average predicted wages for vacancies in other markets	4.8	0.5[b]	21%
Lamadon, Mogstad, and Setzler (2022)	US employer-employee panel data of workers aged 25-60 based on tax filings, 2001-2015	Nested logit model. Instruments (auctions and Bartik) to recover firm level labor supply elasticity.	6.5	4.6[c]	15%
Roussille and Scuderi (2023)	Hired.com data on San Francisco-based job candidates	Very detailed observational data on the hiring process	3.6-5.7		24%[d]
Kroft, Luo, Mogstad, and Setzler (2022)	Employer-employee panel data on US construction firms and procurement auctions, 2001-2015	Estimate inverse labor supply elasticity, with exogenous employment variation arising from firms winning procurement auctions	4.1		25-29%[e]

[a] All estimated markdowns are reported as percentages calculated as (MRPL-w)/w, where MRPL is marginal revenue product of labor and w is the wage.

[b] Market defined as SOC-6 occupation and commuting zone.

[c] Market defined as two-digit industry and commuting zone.

[d] Under the preferred, monopsonistic model, the average predicted markdown is $31,640 (or 19.5% of productivity). So (MRPL-w)/MRPL=1-w/MRPL=0.195. So w/MRPL=0.805, MRPL/w=1.24.

[e] They report markdowns in Table 3. "The first column of Panel A provides our estimate of the wage markdown relative to MRPL, $(1 + \theta)$-1, using either the DiD or RDD estimand." Converting these estimates of .803 and .777 to the (MRPL-w)/w terms gives (1-.803)/.803 and (1-.777)/.777, or 25% and 29%.

3.7 Production function approach

The production function approach offers a distinct method for estimating the markdown "directly" by estimating the marginal revenue product and comparing it with the wage. This contrasts with other methods that rely on the labor supply elasticity. The production function approach can accommodate the possibility that firms do not always set profit-maximizing wages based on the labor supply elasticity, due to factors like adjustment costs and fairness considerations. Consequently, the production function approach emerges as a valuable alternative to estimating the markdown by relying primarily on estimating the marginal revenue product of labor instead of the labor supply elasticity.

Note that the production function literature tends to define markups or markdowns as price over marginal cost, whereas the theoretical industrial organization (IO) literature tends to define it as price minus marginal cost over price (or, for markdowns, marginal product minus wage over wage). Therefore, the reader should be careful when comparing the formulas for the markdown and markup in this section with the formulas in section 2 above, as the latter follow the IO convention.

To understand how this production function method works, consider first the formula for the value of the marginal product of labor L (so the VMPL, not the marginal *revenue* product, see discussion in section 2.3.6) divided by the wage, given a production function $F(L)$:

$$\frac{VMPL}{W} = \frac{P\frac{\partial F}{\partial L}}{W} \tag{78}$$

If we multiply both the numerator and the denominator by the quantity labor share of output L/Q (and we move the price P from the numerator to the inverse of P in the denominator), then we obtain a similar expression, commonly referred to as the "labor wedge". This expression is in terms of the elasticity of output with respect to labor θ^L in the numerator, and the labor share α^L in the denominator.[25] The labor wedge is:

$$\frac{\frac{\partial F}{\partial L}\frac{L}{Q}}{\frac{W}{P}\frac{L}{Q}} = \frac{\theta^L}{\alpha^L} \tag{79}$$

The labor wedge uses VMPL, not MRPL, which makes a difference when firms have market power in both product and labor markets (see Section 2.3.6). When using MRPL, the formula for ν^L, the markdown for the labor input, is:

$$\nu^L \equiv \frac{MRPL}{W} = \left(1 + \frac{P'(F(L))F(L)}{P(F(L))}\right)\frac{P\frac{\partial F}{\partial L}}{W} = \frac{1}{\mu}\frac{\theta^L}{\alpha^L}, \tag{80}$$

where $\mu \equiv P/C'$ is price over marginal cost, i.e. the product market markup as defined in the production function literature.

Therefore, the following holds for every factor of production i:

$$\nu^i \mu = \frac{\theta^i}{\alpha^i} \tag{81}$$

[25] See above equation (42) for a similar formula we derived in the case where market power arises from product and job differentiation.

If there is an input M with no markdown (i.e, $\nu^M = 1$), then we can use that input's wedge to estimate the markup:

$$\mu = \frac{\theta^M}{\alpha^M} \tag{82}$$

Having estimated the markup μ, we can plug it into equation (80), to obtain the markdown for monopsonistic factor L:

$$\nu^L = \frac{\theta^L/\alpha^L}{\theta^M/\alpha^M} \tag{83}$$

Yeh, Macaluso, and Hershbein (2022) use this approach with data from the Census of Manufactures (CM) and the Annual Survey of Manufactures (ASM) for 1976-2014 to estimate markdowns in the US manufacturing sector. They first identify the markup μ as in equation (82) based on the wedge for materials, a non-labor flexible input that is assumed not to be subject to monopsony power. They then combine this markup estimate with the labor elasticity of output and the labor share to obtain the markdown as in equation (83). They find an average markdown of 53% across establishments.[26] The implied average firm-level labor supply elasticity is 1.88.

Mertens (2022) estimates markdowns and markups for manufacturing firms in Germany over the period 1995-2014 using the production function approach. The administrative dataset he uses contains data on both quantities and prices, which allows him to identify both the markup and the markdown (see discussion below page 42 about price data). He finds that, over 1995-2014, there was a small increase in markups (defined as price over marginal cost), from 1 to 1.03 and a large increase in markdowns (defined as marginal revenue product of labor over the wage), from 1.31 to 1.42, over this period. The 1.31-1.42 markdowns correspond to 31%-42% in our preferred definition of the markdown. The study concludes that rising market power can explain about half of the decline in the labor share for the German manufacturing sector, and that almost all of the increase in market power is due to labor rather than product market power.

Brooks, Kaboski, Li, and Qian (2021) use the production function approach to estimate markdowns for manufacturing firms in China and India. To account for econometric issues in the estimation of the markup (see below), they use different estimation approaches, which end up giving similar markdown estimates for Chinese and Indian firms. They find that the employment-weighted average markdown is 16% in China and 18% in India.[27] They also estimate the market-level inverse elasticities of labor supply by regressing the estimated markdowns on the firms' market shares (as in equation (2) from the oligopsony model discussed in section 2.2), and find an elasticity of about 2.2 in both China and India.

A common criticism in the production function estimation literature is that output production functions estimated with revenue data (as opposed to quantity and price data, which are generally not available) are not always well-identified. Bond, Hashemi, Kaplan, and Zoch (2021) argue that the production function approach fails to provide any information on the markup, because when calculated using the revenue elasticity, the markup estimator from Loecker and Warzynski (2012)

[26] The markdown for the average worker is higher because larger firms have higher markdowns (see their Figure 1).

[27] The unweighted average markdowns are 3% in China, and 1% in India. Interestingly, the median markdown is negative in both countries.

is equal to one when there are no input market frictions.[28] However, Yeh, Macaluso, and Hershbein (2022) show that, while this criticism applies to markup estimation, it does not affect markdown estimation, because the biases in the ratio of the labor elasticity and the materials elasticity (see equation 83) cancel each other out.[29]

One of the main assumptions in the production function approach to estimating markdowns is that there is a frictionless input. Hashemi, Kirov, and Traina (2022) propose a direct method for estimating markdowns that bypasses the initial markup calculation, and therefore does not rely on the assumption that there is at least one frictionless input. They propose estimating markdowns based on estimation of the *revenue* production function and the Marginal Revenue Product of Labor (MRPL). For this approach to work, the econometrician needs data on the input's price by firm. In the case of labor, these prices are wages, and they are available in many widely used datasets for production function estimation. This makes labor monopsony estimation a potentially attractive application of this method.

The estimates from the articles discussed in this section utilizing the production function approach are summarized in Table 6.

[28]This is due to the downward bias pointed out by Klette and Griliches (1996). Bond et al. (2021) go a step further and show that this downward bias effectively renders the estimate uninformative about markups. In the presence of input market frictions, the derivation from Bond et al. (2021) shows that the markup estimator recovers the input market frictions (which explains why the literature does not, in fact, find a markup equal to one with this method).

[29]Other researchers have shown that, with more assumptions (such as constant returns to scale) and more sophisticated methods, one can recover the markup using revenue data (Flynn, Traina, and Gandhi, 2019; Kirov and Traina, 2021, 2023). Both Yeh, Macaluso, and Hershbein (2022) and Brooks et al. (2021) implement a version of the constant-returns to scale approach to markup estimation.

Table 6: Elasticity of labor supply and markdowns across studies using mergers and concentration, calibration and simulation, and the production function approach

Study	Data set	Identification	Firm-level elasticity	Market-level elasticity	Estimated markdown[a]
Arnold (2020)	US. Longitudinal Business Database and Longitudinal Employer Household Dynamics (LEHD), 1999-2009	*Mergers and concentration*: Response of wages and employment to quasi-exogenous mergers and acquisitions		0.87[b]	4-5%
Berger, Herkenhoff, and Mongey (2022a)	US Census Longitudinal Business Database, 1977-2011.	*Calibration and simulation*: Calibrate differentiated jobs oligopsony model using reduced form estimates of labor supply elasticity based on variation in state-level corporate tax rates	2.6		39%[c]
Yeh, Macaluso, Hershbein (2022)	US. The Census of Manufactures (CM) and the Annual Survey of Manufactures (ASM), 1976-2014.	*Production function approach*: Estimate the production function to recover the marginal revenue product of labor.	1.88		53%
Mertens (2022)	Administrative data on German manufacturing firms, 1995-2014.	*Production function approach*: Estimate the production function to recover the marginal revenue product of labor.			31%-42%[d]
Brooks, Kaboski, Li, and Qian (2021)	Annual Survey of Chinese Industrial Enterprises, 1999-2007; and India's Annual Survey of Industries, 1999-2011	*Production function approach*: Estimate markdown using production function approach. Regress markdown on market share to recover inverse market-level elasticity.[e]		2.2[f], 2.1[g]	16%[f], 18%[g]

[a] All estimated markdowns are reported as percentages calculated as (MRPL-w)/w, where MRPL is marginal revenue product of labor and w is the wage.
[b] Market defined as four-digit industry and commuting zone.
[c] They report a markdown of w/MRPL=0.72, so this corresponds to 39% markdown in our definition (MRPL-w)/w.
[d] They report the markdown as $MRPL/w$, and we converted it to $MRPL/w - 1 = (MRPL - w)/w$.
[e] Market defined as 4-digit industry province-level for China and state-level for India.
[f] China
[g] India

3.8 Summary and discussion of markdown estimates

Tables 4, 5 and 6 present the summary of firm-level elasticities, market-level elasticities and markdowns. We always report our preferred estimation of the markdown as $(MRPL - w)/w$, adding footnotes when we converted the authors' reported markdown into our preferred format. For market-level elasticities, we report in footnotes how the market is defined.

Firm level elasticities of labor supply range between about 2 and 6, with the exception of a very low elasticity of 0.096 for MTurk workers, and a very high elasticity of 28.6 for non tenure-track faculty (Table 4). Market level elasticities of labor supply are less commonly estimated, and they range between 0.5 and 5.

The markdown is reported in most studies and typically ranges between about 15% and 50%, so that eliminating monopsony power would all other things equal increase wages by 15 to 50%. Equivalently, we can say that wages are equal to between 67% (1/(1+0.5)) and 87% (1/(1+0.15)) of workers' marginal revenue product. There are a few estimates of the markdown that fall outside this range, with a low estimate of 4-5% based on merger effects (Table 6), and a high estimate of 10.4 on MTurk (Table 4).

The fact that the two markdown estimates using the production function approach are not widely different from other estimates is mildly encouraging. Indeed, the production function approach estimates the marginal revenue product of labor, whereas most alternative estimates rely on the labor supply elasticity, alongside the assumption that firms optimize wages to maximize profits in light of the labor supply elasticity. These two kinds of approaches therefore rely on different core assumptions, and it would thus be useful for future research to confirm that the two approaches yield similar results.

Are the estimated markdowns using micro data too large to be reconciled with macro-level statistics such as the labor share, profit share, etc.? First, one needs to be careful when using the micro estimates in a macro context, as some estimates in the literature give more weight to large firms. This can occur either because of data availability or because of estimation constraints when firm fixed effects are used. If large firms have higher markdowns, as found by Yeh, Macaluso, and Hershbein (2022); Brooks et al. (2021), then micro-level markdowns overestimate the average markdown in the economy.

That said, the extent to which micro-level evidence can be reconciled with the aggregate statistics is still an open question. Consider, for example, the labor share, which, according to the BLS, has ranged from between 56% and 64% since 2000. As a back-of-the-envelope calculation, if the aggregate production function is Cobb-Douglas $Y = AK^{\alpha}L^{1-\alpha}$ with an $\alpha = 0.3$, the elasticity of labor supply to the firm is 5 (implying a markdown of 20%), and the product market elasticity is 7 (implying a markup of 14%), then the labor share in the economy is 50%.[30] This is lower than the 56% labor share, but arguably not completely out of the ballpark. Moreover, this is a back-of-the-envelope calculation that has several degrees of freedom and relies to some extent on assumptions about parameters for which there is no consensus in the literature. For example, if the product market elasticity for the aggregate economy were 20 instead of 7, then the labor share would be 55.4%, much closer to the BLS numbers. Further, if the actual aggregate labor supply were higher than 5 (let's say because of underrepresentation of small firms in the existing estimates), the implied labor share would be even higher. We conclude that the question of whether the estimated markdowns are in line with macro aggregates remains unresolved in the literature, indicating that this issue is still very much open for further investigation.

While the empirical results clearly support the existence of monopsony power, one can also glean some insights about the empirical relevance of the different types of monopsony models. Specifically, what do the empirical estimates of the markdown tell us about which mechanisms of monopsony (see Table 3) are empirically relevant? There is empirical support for the oligopsony model and for the job differentiation model, while the empirical support for search is still developing. In the oligopsony model, a finite number of firms engage in strategic interactions and face a finite labor supply elasticity to the market. Evidence about concentration and mergers supports the strategic interaction aspect, while the relatively low estimated labor supply elasticity to the market supports the finite labor supply aspect. In the job differentiation model, jobs are imperfect substitutes. Evidence from nested logit models (Azar, Berry, and Marinescu, 2022; Lamadon, Mogstad, and Setzler, 2022) and calibrated oligopsony models with job differentiation (Berger, Herkenhoff, and Mongey, 2022a) support the idea that jobs are imperfect substitutes. This differentiation also helps explain other empirical facts, such as the existence of wage dispersion. The empirical evidence for search as a mechanism behind monopsony power is less straightforward and more limited in the current state of the literature. Jarosch, Nimczik, and Sorkin (2019) calibrate a model with search, but this approach is not highly informative about whether search is important in explaining the markdown. Some literature finds that large firms have higher markdowns (Yeh, Macaluso, and Hershbein, 2022; Brooks et al., 2021), which is incompatible with simple search models but can be rationalized with more complex search models, e.g. with on the job search (see section 2.5.1).

The literature does not currently allow us to rule out any of the models, but rather shows some support for all of them. This does not mean that every analysis must include oligopsony, job differentiation, and search. Instead, researchers should focus on the mechanism most relevant to their application, while considering whether other mechanisms might confound their results.

Overall, the empirical literature yields increasingly reliable estimates of monopsony power,

[30] Specifically, we calculate the labor share as

$$LS = (1-\alpha) \times \frac{1-\frac{1}{\sigma}}{1+\frac{1}{\theta}} \tag{84}$$

where α is the exponent of capital in the Cobb-Douglas, θ is the elasticity of labor supply to the firm, and σ is the product market elasticity. To derive this equation, we start from the first-order condition of the firm, which is similar to equation (38), but using the marginal product of labor in the Cobb-Douglas production function, and assuming a representative firm so that its price is equal to the price level, and its wage is equal to the wage level in the economy in equilibrium. Operating on that equation implies $\frac{W\times L}{P\times Y} = (1-\alpha)\times\frac{1-\frac{1}{\sigma}}{1+\frac{1}{\theta}}$.

with markdowns typically in the range of 15% to 50%, implying that completely eliminating monopsony power would increase wages by 15 to 50%. However, it is important to recognize that real-world policy interventions would diminish, rather than eliminate, monopsony power, setting these figures as the theoretical maximum for wage increases achievable solely through monopsony power reduction. Nonetheless, the estimated markdowns indicate substantial scope for meticulously crafted policies to bolster wages by curtailing monopsony power.

4 Policy and monopsony power

In this section, we focus on three policies that have been carefully studied in connection with monopsony power: merger control in antitrust policy, non-competition agreements, and the minimum wage. We discussed these policies and other policies related to monopsony power in a policy-focused literature review (Azar and Marinescu, forthcoming).

4.1 Merger control

Antitrust authorities have the power to sue to block anticompetitive mergers. Merger review is fundamentally a prediction exercise, where antitrust authorities assess the risk that a merger will generate anticompetitive effects.

A merger of competing employers can substantially increase labor market power. Indeed, a merger increases the market share of the merging parties. For this reason, a merger is predicted to increase the markdown either in an oligopsony model (potentially with job differentiation, as in Berger, Herkenhoff, and Mongey (2022a) or Azar, Berry, and Marinescu (2022)), or in a search and matching model that uses the HHI as a measure of market power (Jarosch, Nimczik, and Sorkin, 2019). While an increase in productivity from the merger (commonly referred to as "efficiencies") can theoretically offset the increase in market power (see section 3.3.1), evidence shows that mergers that substantially increase labor market concentration reduce wages (Prager and Schmitt, 2021; Arnold, 2019). This suggests that efficiencies are, on average, not large enough to lead to an increase in wages or a decrease in the markdown.

The remainder of this section focuses on American antitrust policy, but similar approaches are used in merger enforcement by competition authorities around the world. Until recently, US antitrust enforcement focused on competition in the product market, and thus mergers affecting the labor market received little to no attention (Marinescu and Posner, 2019; Marinescu and Hovenkamp, 2019). In a significant development, the US Department of Justice Antitrust Division successfully sued to block a publisher merger that would have reduced the compensation for labor. The merger between two major book publishers, Penguin-Random House and Simon & Schuster, was blocked in October 2022. The Antitrust Division predicted that authors would see lower pay as a result of the merger, because the merger would lead to reduced competition in bidding for authors' manuscripts. This was the first merger blocked in the US on the grounds of labor monopsony.

The US antitrust authorities – the Federal Trade Commission (FTC) and the Department of Justice Antitrust Division – periodically publish merger guidelines. The guidelines describe how agencies determine which mergers they will sue to block. They are designed to help the public, business community, practitioners, and courts understand the factors and frameworks the Agencies consider when investigating mergers (Federal Trade Commission, 2022).

While the merger guidelines mentioned the issue of buyer power as early as 1982 (in footnote 5), there was at first no explicit discussion of the power of employers as buyers of labor services. The 2010 guidelines expanded the discussion of monopsony with a whole section on powerful buyers. Finally, the 2023 Merger Guidelines (U.S. Department of Justice and the Federal Trade Commission, 2023) included Guideline 10, which is dedicated to the analysis of mergers that may substantially lessen competition for workers, creators, suppliers or other providers.

There are several tools that can be used to assess the effects of mergers on labor market power. In guidelines 1 through 6, the 2023 merger guidelines identify six distinct frameworks that can be used to assess whether a merger may substantially lessen competition.

Guideline 1 focuses on concentration, which is a key indicator that can be used by the antitrust authorities in all merger reviews, whether involving the labor market or other markets. A merger that significantly increases concentration in a highly concentrated market will attract the attention of US antitrust authorities (see Guideline 1 of the 2023 Merger Guidelines (U.S. Department of Justice and the Federal Trade Commission, 2023)). Concentration is a useful indicator for a merger's likely effects on competition: a merger reduces direct competition between the merging firms (also see Guideline 2), and typically facilitates coordination (i.e. tacit or explicit collusion; covered in Guideline 3) among the remaining firms (Baker and Farrell, 2020).

For the product market, there is wide-ranging empirical evidence that mergers that increase concentration increase prices on average (e.g. Kwoka, 2014; Bhattacharya, Illanes, and Stillerman, 2023). For the labor market, we have mentioned parallel results above: mergers that increase labor market concentration tend to decrease wages (Prager and Schmitt, 2021; Arnold, 2019).

From a legal standpoint, a merger that significantly increases the HHI in a highly concentrated market leads to a "structural presumption". The structural presumption shifts the burden of proof to the merging parties: unless sufficient contradictory evidence or rebuttal is provided by the merging firms, the antitrust authorities *presume* that such a merger is anticompetitive.[31] Specifically, markets with an HHI greater than 1,800 are highly concentrated, and a merger in a highly concentrated market that increases the HHI by more than 100 points is presumed to substantially lessen competition or tend to create a monopoly (U.S. Department of Justice and the Federal Trade Commission, 2023).[32]

In order to calculate the HHI, it is necessary to define a market where *market shares* can be calculated. There is no one single market that is "correct". Too broad a market would make sure that the market captures all relevant substitutes, but would include many options that are not reasonable substitutes. Too narrow a market would only include the very closest substitutes but would miss relevant competition from slightly more distant substitutes. Recognizing this continuum, the antitrust authorities do not define "the market" but "a relevant market" where the loss of competition stemming from a proposed merger can be well understood.

[31] Rebuttal evidence can be provided by showing that the acquired firm is failing, that there will be firm entry or product repositioning, or that the merger brings efficiencies, i.e. higher productivity. In all cases, the merging parties must demonstrate that, based on the rebuttal evidence, the merger is not threatening a substantial reduction in competition (U.S. Department of Justice and the Federal Trade Commission, 2023).

[32] The HHI thresholds have been revised downward in the 2023 guidelines relative to the 2010 guidelines, going back to the thresholds that were used before 2010, and that were established by the 1982 merger guidelines.

The 2023 Merger Guidelines offer tools to define a relevant *labor* antitrust market in Section 4.3.D.8. Broadly, a labor market is defined so that jobs are reasonably substitutable from the point of view of workers: "[d]epending on the occupation, alternative job opportunities might include the same occupation with alternative employers, or alternative occupations. Geographic market definition may involve considering workers' willingness or ability to commute, including the availability of public transportation." One tool to define a labor market based on this substitutability principle is the hypothetical monopsonist test: the test asks whether it would be profitable for a hypothetical monopsonist to undertake at least a small but significant (a 5% threshold is often considered) and non-transitory decrease in wage or other worsening of terms or working conditions. When considering a change in wages or other terms, current market conditions are often taken as the benchmark, but other more competitive benchmarks can be used.[33] When the market level labor supply is low enough, workers are unable to sufficiently substitute out of the market to defeat the monopsonist's strategy. In this case, the monopsonist finds it profitable to impose lower wages or worse terms, and the candidate labor market passes the test, and can be thought of as a relevant antitrust market.[34]

While we have focused so far on the role of concentration and market definition in merger review, there are five additional core frameworks that agencies use to analyze mergers (Guidelines 2-6). Here, we briefly highlight how a merger affecting the labor market can be analyzed under Guidelines 2 or 3. Guideline 2 describes how a merger leads to the loss of head-to-head competition between the merging parties. In the case of the labor market, such "head-to-head" competition can be measured, for example, by the workers' propensity to move from one of the merging firms to the other in response to changes in wage or working conditions. Guideline 3 describes the risk of coordination: after a merger, there are fewer employers in a labor market, which can make it easier for the remaining employers to coordinate, e.g. by agreeing on wages, or agreeing not to solicit each other's employees.

To sum up, antitrust authorities have paid increasing attention to mergers that may reduce competition for labor. The US Department of Justice won its first case in this area in 2022, and the 2023 Merger Guidelines explain that anticompetitive effects in the labor market alone are sufficient for the antitrust authorities to sue to block a merger. All along, research in economics has informed this policy agenda.

4.2 Non-competition agreements

Non-compete clauses are legal agreements that prevent employees from joining competing firms within a specified time frame after their employment ends. Theoretically, non-competes can serve as a response to the employer's holdup problem when investing in human capital (Hart and Moore, 1990). Employers may hesitate to provide extensive training or share sensitive information if employees can subsequently transfer this knowledge to a rival firm (Barron, Berger, and Black, 1999; Acemoglu and Pischke, 1999). Enforceable non-compete agreements mitigate this risk by

[33] When a market is already close to a monopsony, a hypothetical monopsonist may not find it profitable to impose a *further* 5% decrease in wages because the dominant firm has already used most of the market power of a monopsonist. This means that, paradoxically, monopsony labor markets would fail the test. This is why the agencies can decide to use more competitive benchmarks than the current market conditions. This paradoxical situation is referred to in antitrust jargon as the *cellophane fallacy*.

[34] Note that there can be more than one candidate market that passes this test.

ensuring that employees cannot immediately benefit a competitor with the training and insights they have gained (Rubin and Shedd, 1981; Posner, Triantis, and Triantis, 2004). However, while non-competes increase incentives to provide general training, they can reduce incentives to provide firm-specific training (Meccheri, 2009). This is because under non-competes, employers are not worried about losing their workers to competitors and therefore can provide them with general training and skills applicable to all companies rather than just firm-specific training.

Although non-competes can safeguard the employer's investment, they also influence the dynamics of labor market competition. These agreements can suppress wages and limit worker mobility (Krueger and Posner, 2018; Marx, 2018). When viewed through the search model framework, non-competes effectively narrow the pool of potential employers, thereby increasing market concentration and reducing the bargaining power of employees (Jarosch, Nimczik, and Sorkin, 2019). Viewed through an oligopsony model, a firm with a non-compete becomes a monopsonist for its current employees. However, the impact of non-compete agreements varies depending on whether an individual is currently bound by such an agreement or is a prospective employee without existing restrictions. The former may face limited labor market options, while the latter encounters a more competitive environment. Nonetheless, wage equality constraints within firms – as in the search model by Rudanko (2023) (or in the oligopsony model) – can lead to adverse wage outcomes for all workers, regardless of their non-compete status.

Empirical research consistently shows that non-compete agreements tend to suppress wages. A significant focus of the empirical literature has been on the enforceability of non-competes. The term "enforceability" in this context means the likelihood of a court upholding a non-compete's restrictions on an employee's post-employment activities. The degree to which non-competes are enforceable varies by state law and individual case specifics. For example, California has outlawed non-competes, while other states have set certain restrictions or conditions for their application. Research by Starr (2019) indicates that stronger enforcement of non-competes led to increases in employer-provided training but also reduced wages by 4%.

Non-compete agreements are surprisingly widespread across various job sectors in the United States. Starr, Prescott, and Bishara (2021) reveal that 18.1% of U.S. workers are subject to non-competes, with these agreements being more common in high-skill occupations but also widespread in low-skill occupations. Only a small percentage of employees negotiate the terms of non-competes. The same research found that stricter enforcement is associated with 1.7% to 2.5% lower wages. However, non-competes offered before a job offer[35] tend to result in better wage outcomes, suggesting that workers may receive compensation for agreeing to these terms.

Non-compete agreements can disproportionately affect wages for certain demographic groups, and their impact can extend beyond individual workers to the labor market as a whole. Johnson, Lavetti, and Lipsitz (2020) explored the impact of changes in state-level enforceability of non-competes, showing that such agreements disproportionately lower wages for women and non-white workers. The study suggests that enforceability curtails workers' leverage to secure higher pay during favorable labor market conditions. Moreover, Johnson, Lavetti, and Lipsitz (2020) indicate that strict non-compete enforcement in one state adversely affects earnings even for workers in other states within the same job market (for local labor markets that straddle state borders). This negative externality challenges the notion that workers consent to non-competes because they are

[35] 61% of non-competes are presented before job offer and 29% after, with the remainder being for promotions/raises, and respondents who couldn't remember the timing.

mutually beneficial contracts.

Banning non-compete agreements can boost worker mobility and wages. Balasubramanian, Chang, Sakakibara, Sivadasan, and Starr (2022) investigated the consequences of Hawaii's non-compete ban on tech employees, finding a 4% wage increase and an 11% boost in job-switching activity. Overall, tech workers in Hawaii saw a 4.6% cumulative earnings gain compared to those in states with moderate non-compete enforceability levels.

Non-compete agreements, while designed to protect business interests, can inadvertently stifle industry growth and innovation. Economic theory suggests that, in sectors where the high benefits of investments are less likely to be passed on to consumers, non-competes may prevent the emergence of spin-offs, to the detriment of consumers (Lipsitz and Tremblay, 2021). A policy analysis (Shi, 2023) extending the Postel-Vinay and Robin (2002) search and matching framework shows that limiting the duration of non-competes for top executives to roughly one month strikes an optimal balance. The calibration show this one-month limit is optimal, as it minimizes the negative impact on companies unable to recruit talent, while still encouraging investment by a worker's current employer. The research suggests that a quasi-ban on non-competes could be optimal, even for employees like executives, who are well-informed and who may indeed have access to knowledge that companies would seek to protect from competitors.

The empirical literature shows that non-competes can depress wages, with low-wage workers being disproportionately affected. These agreements shift the power balance in favor of employers, making it more challenging for employees to negotiate better pay. For low-wage positions, the justification for protecting training investments and proprietary information is weaker. Non-competes can further restrict career advancement and wage increases. Employers, particularly those hiring specialized or high-level staff, can often safeguard their interests post-employment through less restrictive means, such as confidentiality agreements and intellectual property laws. Given the availability of alternative protective measures, it appears that the typical employer places little importance on the enforceability of non-compete clauses (Hiraiwa, Lipsitz, and Starr, 2023).

In a significant policy development, the US Federal Trade Commission (FTC) issued a rule banning non competition agreements (with few exceptions); the rule is set to take effect in September 2024. According to the FTC, non-compete clauses constitute an unfair competitive practice under Section 5 of the FTC Act. In preparation for the rule, the FTC conduced a comprehensive review of the empirical evidence on the effects of non-competes (Federal Trade Commission, 2023). This move indicates the FTC's efforts to address concerns related to employer monopsony power, and to promote more competitive labor markets.

4.3 Minimum wage

Economic theory predicts that competitive labor markets typically see a decrease in employment when a minimum wage is implemented. However, as explained above, employers with monopsony power pay workers less than the marginal revenue product of labor. Therefore, employers could afford to pay workers somewhat more. This implies that the effect of the minimum wage on employment is ambiguous (Robinson, [1933] 1966; Manning, 2013, 2021): it depends on the degree of competition in the labor market, and on the level of the minimum wage relative to the marginal revenue product of labor.

To examine the employment effect of the minimum wage, we can use different models: an oligopsony framework (with or without job differentiation), or a random search model with posting and on the job search (Burdett and Mortensen, 1998). As explained in the theory section 2 above, under oligopsony, both wages and employment levels are suppressed compared to a perfectly competitive market. Firms aiming to maximize profits will set wages at a point where the additional cost of hiring one more employee (the wage of the new employee, plus the increase in the wage bill for all incumbents) equals the marginal revenue product. Introducing a minimum wage changes the optimal choice for the firm. When the minimum wage is set above the monopsonistic equilibrium but does not exceed the competitive minimum wage, the marginal cost of hiring an additional employee is simply the minimum wage. Firms can thus hire more workers without having to raise pay for existing employees, as they must already earn at least the minimum wage. This reduced marginal cost allows the monopsonist to hire more workers at the minimum wage, up to the point where the last worker's marginal revenue product equals the minimum wage (Bhaskar, Manning, and To, 2002). To summarize, when the minimum wage increases while staying below the competitive level, employment increases in a monopsonistic labor market.

Azar et al. (2023) show that the negative impact of a local minimum wage on employment is mitigated – or can even turn positive – in more concentrated labor markets. This outcome can be predicted by a simple oligopsony model with Cournot competition, where the only difference between labor markets is the number of employers. Markets with fewer employers have a higher HHI and lower wages. When the minimum wage is set relative to the equilibrium wage prevailing before the policy change, the employment impact is more positive in markets with higher labor market concentration. This is because in more concentrated markets, markups are higher, leaving more scope to raise wages without cutting employment. Empirical evidence from local minimum wage hikes in the US align with these predictions, suggesting that monopsony power may help account for the varying effects of minimum wage on employment found in the literature.

Minimum wage effects can also be modeled in a differentiated jobs framework, as in Dustmann et al. (2022). Dustmann et al. (2022) study the introduction of the national minimum wage in Germany in 2015. They assess the policy's impact using administrative data and quasi-exogenous variation in exposure across different German labor markets. They use high-wage workers as a control group for low-wage workers who were affected by the minimum wage. They find that the introduction of the minimum wage increased wages but did not decrease employment. The policy instead reallocated workers toward higher paying firms that are also larger and more desirable as measured by the poaching rank (i.e. how likely the firm is to be chosen over other firms among workers who have a job-to-job transition). The reallocation effect explains 17% of the wage increase induced by the introduction of the minimum wage. The differentiated jobs model – which is similar to the model in section 2.3.3 – explains the reallocation effect in a monopsony framework with heterogeneous firms: lower wage and less productive firms are unable to pay workers the new minimum wage and lose employment, while higher wage firms are induced to pay more, and therefore gain employment. The model predicts ambiguous wage and employment effects on the highest wage firms. Despite these ambiguous theoretical effects for the most productive firms, average worker welfare increases in the model as long as the minimum wage employment effect is not negative, which is the case that is empirically relevant in Dustmann et al. (2022).

Macroeconomic models suggest that minimum wages can increase social welfare through both efficiency gains and redistribution. Berger, Herkenhoff, and Mongey (2022b) calibrate a general equilibrium oligopsony model. The calibration indicates that the welfare-maximizing minimum wage is \$15 per hour, with a range from \$0 to \$31 depending on the social welfare function weights. However, focusing solely on efficiency and ignoring benefits from redistribution, the optimal minimum wage drops to approximately \$8 per hour, and generates only marginal gains in efficiency.

In Japan, just as in the US, the extent of monopsony power across labor markets has been quantified using a production function approach. Okudaira, Takizawa, and Yamanouchi (2019) find a negative impact of minimum wages on employment in the more competitive labor markets where minimum wages approach the marginal product of labor. However, in markets with significant monopsony power, minimum wage increases do not affect employment levels.

A study examining Walmart's wage structure found that workers place substantial value on non-wage amenities. This preference is evident in workers' willingness to pay for "dignity at work," as measured using survey experiments (Dube, Naidu, and Reich, 2022). Following Walmart's corporate minimum wage hike in 2014, there was no observed reduction in job amenities, suggesting that higher minimum wages may enhance the overall job value without necessitating a cutback in other benefits. This contradicts the traditional compensating differentials theory, which would predict a decrease in non-wage amenities in response to wage increases. The findings align more closely with a monopsonistic labor market perspective based on job differentiation, where wages and benefits are not highly substitutable from the employees' viewpoint. [36]

The minimum wage has historically been a powerful tool for narrowing the earnings gap between black and white workers, particularly during the pivotal period of the late 1960s and early 1970s in the United States. However, if, as predicted by the competitive theory of the labor market, the minimum wage had disproportionately reduced black employment, the minimum wage's effect on overall racial inequality in the labor market would have been dampened. Derenoncourt and Montialoux (2021) examine the impact of the expansion of federal minimum wage laws (Fair Labor Standards Act) in 1966 to include industries with high shares of black workers, such as agriculture, restaurants, and personal services. They find that the racial earnings gap, when adjusted for measurable factors, was virtually eliminated within these industries post-reform, thus falling by 25 log points. Moreover, the research indicates that this policy did not adversely affect the employment of black workers, nor did it have a negative impact on overall employment levels. The findings are therefore consistent with the existence of monopsony power in the labor market, which allowed employers to increase wages without cutting employment.

This evidence underscores the potential of the minimum wage to serve as an effective mechanism for boosting earnings for the lowest-paid workers. The absence of pervasive negative employment consequences suggests that employers' monopsony power may play a role in these outcomes. In a purely competitive market, raising the minimum wage is predicted to lower employment. However, in labor markets with substantial monopsony power, an increase in the minimum wage may simultaneously boost employment and wages.

[36] However, Walmart's self-imposed wage increase may stem from a strategic business decision to boost overall compensation. Such a decision would not necessarily lead to a reduction in amenities, and firms' reaction to a minimum wage imposed by regulation may be different.

5 Conclusion

This chapter provides a comprehensive theoretical and empirical examination of monopsony power in labor markets. Our chapter summarizes recent theoretical contributions from three classes of models: oligopsony, job differentiation, and search and matching. The markdown is our preferred summary statistic for monopsony power: we define it as $(MRPL - w)/w$, the percent wage increase that would occur if monopsony power were eliminated, so that wages equal the marginal revenue product of labor ($w = MRPL$). We derive the markdown for each theoretical framework and show that the average markdown often depends on labor market concentration, with higher concentration typically leading to higher markdowns in equilibrium. Beyond concentration, we also identify other determinants of market power, such as job differentiation and the labor supply elasticity, which are crucial for a comprehensive understanding of monopsony power.

The empirical literature has used different approaches to quantifying labor market power, using labor market concentration, the labor supply elasticity, and the production function approach. A first strand of the literature estimates the labor supply elasticity using exogenous or quasi-exogenous variation, and often proceeds to infer the markdown as the inverse of the labor supply elasticity. A second strand estimates the relationship between labor market concentration and wages, and must contend with productivity as a key potential confound; this literature typically finds that labor market concentration lowers wages. A third strand uses calibration and simulation to infer the markdown. And a fourth strand estimates the markdown using structural estimation or a production function approach. Across the overwhelming majority of studies, markdown estimates range from 15% to 50%. Most labor supply elasticity estimates to the firm fall between 2 and 6, and estimates of the labor supply elasticity to the market fall between 0.5 and 5, though market-level elasticities naturally depend on the definition of the market.

Given these substantial markdowns, we conclude that wages are not exclusively determined by marginal productivity. Therefore, variations in the markdown should also be considered as a source of wage variation, and researchers should be careful not to interpret the wage as a fully reliable measure of the marginal product of labor.

We then investigate policy implications of monopsony power, discussing merger control, non-competition agreements and the minimum wage. The literature suggests that mergers can lead to anticompetitive effects in labor markets similar to those in product markets. Monopsony power also helps reconcile the disparate estimates of the employment effects of the minimum wage: when there is enough monopsony power, a minimum wage increase may not reduce employment, and could in fact increase it.

While the literature has made strides in understanding monopsony power, there are certain limitations that future research could helpfully address. We outline three main limitations here. First, the treatment of non-wage amenities in the presence of monopsony power could be improved. The literature has generally assumed amenities are fixed and costless to firms; it would be interesting to examine how key conclusions are affected when detailed data on amenities is used, and when the endogeneity of firms' amenity choices is more fully addressed. Second, the empirical literature estimating the markdown could better integrate the estimation of labor and product market power, seeking a robust estimation of key parameters. Third, future research could help us more systematically understand which sources of monopsony power – oligopsony, job differentiation, and search and matching frictions – are most quantitatively important in different contexts. This exploration would be helpful to determine which policy interventions are most likely to succeed in addressing the adverse effects of monopsony power.

Overall, the literature has progressed in a few ways since Manning's chapter in this same Handbook (Manning, 2011). First, it has developed new models of monopsony power, drawing both on the IO literature, and on labor economics' search and matching models. Second, the estimates of the markdown are getting more reliable. And third, the proliferation of new theories and empirical estimates has been accompanied by new insights on policy, showing that mergers and non-competition agreements may negatively affect wages by reducing competition for labor, and that moderate minimum wages may benefit workers by reducing the markdown and increasing employment. In essence, the recent surge in monopsony research has catalyzed a critical reevaluation of our understanding of labor markets, propelling the once-marginalized notion of imperfect competition to the forefront of economic analysis and policy.

References

Abowd, John M., Francis Kramarz, and David N. Margolis. 1999. "High Wage Workers and High Wage Firms." *Econometrica* 67 (2):251–333. URL http://www.jstor.org/stable/2999586.

Acemoglu, Daron and Jörn-Steffen Pischke. 1999. "The Structure of Wages and Investment in General Training." *Journal of Political Economy* 107 (3):539–572. URL https://www.journals.uchicago.edu/doi/10.1086/250071.

Amodio, Francesco and Nicolás de Roux. 2023. "Measuring Labor Market Power in Developing Countries: Evidence from Colombian Plants." *Journal of Labor Economics* URL https://www.journals.uchicago.edu/doi/10.1086/725248. Publisher: The University of Chicago Press.

Anderson, Simon P, André De Palma, and Jacques-François Thisse. 1988. "The CES and the logit: Two related models of heterogeneity." *Regional Science and Urban Economics* 18 (1):155–164.

Arnold, David. 2019. "Mergers and Acquisitions, Local Labor Market Concentration, and Worker Outcomes." *SSRN Electronic Journal* URL https://www.ssrn.com/abstract=3476369.

Azar, José. 2012. *A new look at oligopoly: Implicit collusion through portfolio diversification*. Ph.D. thesis, Princeton University.

Azar, José, Steven Berry, and Ioana Elena Marinescu. 2019. "Estimating Labor Market Power." SSRN Scholarly Paper ID 3456277, Social Science Research Network, Rochester, NY. URL https://papers.ssrn.com/abstract=3456277.

Azar, José, Ioana Marinescu, Marshall Steinbaum, and Bledi Taska. 2020. "Concentration in US labor markets: Evidence from online vacancy data." *Labour Economics* 66:101886. URL https://www.sciencedirect.com/science/article/pii/S0927537120300907.

Azar, José, Yue Qiu, and Aaron Sojourner. 2021. "Common Ownership Reduces Wages and Employment." SSRN Scholarly Paper ID 3954399, Social Science Research Network, Rochester, NY. URL https://papers.ssrn.com/abstract=3954399.

———. 2022. "Common ownership in labor markets." *Available at SSRN 4158482* .

Azar, José and Ricardo M Ribeiro. 2022. "Estimating oligopoly with shareholder voting models." .

Azar, José and Xavier Vives. 2021. "General equilibrium oligopoly and ownership structure." *Econometrica* 89 (3):999–1048.

Azar, José, Emiliano Huet-Vaughn, Ioana Marinescu, Bledi Taska, and Till von Wachter. 2023. "Minimum Wage Employment Effects and Labor Market Concentration." *The Review of Economic Studies* URL https://doi.org/10.1093/restud/rdad091.

Azar, José and Ioana Marinescu. forthcoming. "Monopsony Power in the Labor Market: From Theory to Policy." *Annual Review of Economics* 16.

Azar, José A., Steven T. Berry, and Ioana Marinescu. 2022. "Estimating Labor Market Power." Working Paper. URL https://www.nber.org/papers/w30365.

Backus, Matthew, Christopher Conlon, and Michael Sinkinson. 2021. "Common ownership in America: 1980–2017." *American Economic Journal: Microeconomics* 13 (3):273–308.

Baker, Jonathan and Joseph Farrell. 2020. "Oligopoly Coordination, Economic Analysis, and the Prophylactic Role of Horizontal Merger Enforcement." *American University Washington College of Law Research Paper No. 2020-23* URL https://digitalcommons.wcl.american.edu/facsch_lawrev/1916.

Balasubramanian, Natarajan, Jin Woo Chang, Mariko Sakakibara, Jagadeesh Sivadasan, and Evan Starr. 2022. "Locked In? The Enforceability of Covenants Not to Compete and the Careers of High-Tech Workers." *Journal of Human Resources* 57 (S):S349–S396. URL https://jhr.uwpress.org/content/57/S/S349. Publisher: University of Wisconsin Press Section: Articles.

Barron, John M., Mark C. Berger, and Dan A. Black. 1999. "Do Workers Pay for On-The-Job Training?" *The Journal of Human Resources* 34 (2):235. URL https://www.jstor.org/stable/146344?origin=crossref.

Bassier, Ihsaan, Arindrajit Dube, and Suresh Naidu. 2022. "Monopsony in Movers: The Elasticity of Labor Supply to Firm Wage Policies." *Journal of Human Resources* 57 (S):S50–s86. URL http://jhr.uwpress.org/lookup/doi/10.3368/jhr.monopsony.0319-10111R1.

Berger, David, Kyle Herkenhoff, and Simon Mongey. 2019. "Labor Market Power." Tech. Rep. w25719, National Bureau of Economic Research, Cambridge, MA. URL http://www.nber.org/papers/w25719.pdf.

———. 2022a. "Labor Market Power." *American Economic Review* 112 (4):1147–1193. URL https://www.aeaweb.org/articles?id=10.1257/aer.20191521.

Berger, David W., Kyle F. Herkenhoff, and Simon Mongey. 2022b. "Minimum Wages, Efficiency and Welfare." Working Paper. URL https://www.nber.org/papers/w29662.

Bertrand, Joseph. 1883. "Book review of theorie mathematique de la richesse sociale and of recherches sur les principles mathematiques de la theorie des richesses." *Journal de Savants* (67):499–508.

Bhaskar, V, Alan Manning, and Ted To. 2002. "Oligopsony and Monopsonistic Competition in Labor Markets." *Journal of Economic Perspectives* 16 (2):155–174. URL https://pubs.aeaweb.org/doi/10.1257/0895330027300.

Bhattacharya, Vivek, Gastón Illanes, and David Stillerman. 2023. "Merger Effects and Antitrust Enforcement: Evidence from US Consumer Packaged Goods." URL https://www.nber.org/papers/w31123.

Boal, William M. and Michael R Ransom. 1997. "Monopsony in the Labor Market." *Journal of Economic Literature* 35 (1):86–112. URL http://www.jstor.org/stable/2729694.

Boar, Corina and Virgiliu Midrigan. 2024. "Markups and Inequality." *Review of Economic Studies* .

Bond, Steve, Arshia Hashemi, Greg Kaplan, and Piotr Zoch. 2021. "Some unpleasant markup arithmetic: Production function elasticities and their estimation from production data." *Journal of Monetary Economics* 121:1–14.

Bresnahan, Timothy F and Steven C Salop. 1986. "Quantifying the competitive effects of production joint ventures." *International Journal of Industrial Organization* 4 (2):155–175.

Brooks, Wyatt J., Joseph P. Kaboski, Yao Amber Li, and Wei Qian. 2021. "Exploitation of labor? Classical monopsony power and labor's share." *Journal of Development Economics* 150:102627. URL https://www.sciencedirect.com/science/article/pii/S0304387821000043.

Burdett, Kenneth and Dale T. Mortensen. 1998. "Wage Differentials, Employer Size, and Unemployment." *International Economic Review* 39 (2):257–273. URL https://www.jstor.org/stable/2527292. Publisher: [Economics Department of the University of Pennsylvania, Wiley, Institute of Social and Economic Research, Osaka University].

Caldwell, Sydnee and Oren Danieli. 2023. "Outside Options in the Labor Market." *Review of Economic Studies* .

Caldwell, Sydnee and Emily Oehlsen. 2022. "Gender Differences in Labor Supply: Experimental Evidence from the Gig Economy." .

Card, David. 2022. "Who Set Your Wage?" *American Economic Review* 112 (4):1075–1090. URL https://pubs.aeaweb.org/doi/10.1257/aer.112.4.1075.

Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline. 2018. "Firms and Labor Market Inequality: Evidence and Some Theory." *Journal of Labor Economics* 36 (S1):S13–S70. URL https://www.journals.uchicago.edu/doi/10.1086/694153. Publisher: The University of Chicago Press.

Choi, Hyeri and Ioana Marinescu. 2024. "Data for labor market concentration using Lightcast (formerly Burning Glass Technologies)." *Data in Brief* 55:110647. URL https://linkinghub.elsevier.com/retrieve/pii/S2352340924006140.

Cowgill, Bo, Andrea Prat, and Tommaso Valletti. 2023. "Political Power and Market Power." URL http://arxiv.org/abs/2106.13612. ArXiv:2106.13612 [econ, q-fin].

Dal Bó, Ernesto, Frederico Finan, and Martín A. Rossi. 2013. "Strengthening State Capabilities: The Role of Financial Incentives in the Call to Public Service*." *The Quarterly Journal of Economics* 128 (3):1169–1218. URL https://academic.oup.com/qje/article/128/3/1169/1849634.

Davis, Steven J., R. Jason Faberman, and John C. Haltiwanger. 2012. "Recruiting Intensity during and after the Great Recession: National and Industry Evidence." *American Economic Review* 102 (3):584–588. URL https://www.aeaweb.org/articles?id=10.1257%2Faer.102.3.584&ref=employ-america.

Deb, Shubhdeep, Jan Eeckhout, Aseem Patel, and Lawrence Warren. 2022. "What Drives Wage Stagnation: Monopsony or Monopoly?" *Journal of the European Economic Association* 20 (6):2181–2225. URL https://doi.org/10.1093/jeea/jvac060.

Derenoncourt, Ellora and Claire Montialoux. 2021. "Minimum Wages and Racial Inequality*." *The Quarterly Journal of Economics* 136 (1):169–228. URL https://doi.org/10.1093/qje/qjaa031.

Dube, Arindrajit, Jeff Jacobs, Suresh Naidu, and Siddharth Suri. 2020. "Monopsony in Online Labor Markets." *American Economic Review: Insights* 2 (1):33–46. URL https://www.aeaweb.org/articles?id=10.1257/aeri.20180150.

Dube, Arindrajit, Suresh Naidu, and Adam D Reich. 2022. "Power and Dignity in the Low-Wage Labor Market: Theory and Evidence from Wal-Mart Workers." Working Paper 30441, National Bureau of Economic Research. URL http://www.nber.org/papers/w30441.

Dustmann, Christian, Attila Lindner, Uta Schönberg, Matthias Umkehrer, and Philipp vom Berge. 2022. "Reallocation Effects of the Minimum Wage." *The Quarterly Journal of Economics* 137 (1):267–328. URL https://doi.org/10.1093/qje/qjab028.

Federal Trade Commission. 2022. "Federal Trade Commission and Justice Department Seek to Strengthen Enforcement Against Illegal Mergers." URL https://www.ftc.gov/news-events/news/press-releases/2022/01/federal-trade-commission-justice-department-seek-strengthen-enforcement-against-illegal-

———. 2023. "Non-Compete Clause Rule NPRM." Federal Register Notice 88 Fed. Reg. 3482, Federal Trade Commission. URL https://www.federalregister.gov/documents/2023/01/19/2023-00414/non-compete-clause-rule.

Flynn, Zach, James Traina, and Amit Gandhi. 2019. "Measuring markups with production data." *Available at SSRN 3358472* .

Frank, Robert H. 2005. "Positional Externalities Cause Large and Preventable Welfare Losses." *The American Economic Review* 95 (2):137–141. URL https://www.jstor.org/stable/4132805. Publisher: American Economic Association.

Gabszewicz, Jean Jaskold and Jean-Philippe Vial. 1972. "Oligopoly "a la Cournot" in a general equilibrium analysis." *Journal of economic theory* 4 (3):381–400.

Goolsbee, Austan and Chad Syverson. 2023. "Monopsony Power in Higher Education: A Tale of Two Tracks." *Journal of Labor Economics* 41 (S1):S257–S290. URL https://www.journals.uchicago.edu/doi/abs/10.1086/726720. Publisher: The University of Chicago Press.

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

Harford, Jarrad, Dirk Jenter, and Kai Li. 2011. "Institutional cross-holdings and their effect on acquisition decisions." *Journal of Financial Economics* 99 (1):27–39.

Hart, Oliver and John Moore. 1990. "Property Rights and the Nature of the Firm." *Journal of political economy* 98 (6):1119–1158.

Hashemi, Arshia, Ivan Kirov, and James Traina. 2022. "The production approach to markup estimation often measures input distortions." *Economics Letters* 217:110673.

Hiraiwa, Takuya, Michael Lipsitz, and Evan Starr. 2023. "Do Firms Value Court Enforceability of Noncompete Agreements? A Revealed Preference Approach." URL https://papers.ssrn.com/abstract=4364674.

Jarosch, Gregor, Jan Sebastian Nimczik, and Isaac Sorkin. 2019. "Granular Search, Market Structure, and Wages." Tech. Rep. w26239, National Bureau of Economic Research. URL https://www.nber.org/papers/w26239.

Johnson, Matthew S., Kurt Lavetti, and Michael Lipsitz. 2020. "The Labor Market Effects of Legal Restrictions on Worker Mobility." URL https://papers.ssrn.com/abstract=3455381.

Khan, Lina M and Sandeep Vaheesan. 2017. "Market Power and Inequality: The Antitrust Counterrevolution and Its Discontents." *Policy Review* 11.

Kirov, Ivan and James Traina. 2021. "Measuring markups with revenue data." *Available at SSRN 3912966* .

———. 2023. "Labor market power and technological change in US manufacturing." Tech. rep., Working Paper.

Klette, Tor Jakob and Zvi Griliches. 1996. "The inconsistency of common scale estimators when output prices are unobserved and endogenous." *Journal of applied econometrics* 11 (4):343–361.

Kroft, Kory, Yao Luo, Magne Mogstad, and Bradley Setzler. 2022. "Imperfect Competition and Rents in Labor and Product Markets: The Case of the Construction Industry."

Krueger, Alan B and Eric A Posner. 2018. "A Proposal for Protecting Low-Income Workers from Monopsony and Collusion." Policy Proposal 2018-05, The Hamilton Project. URL https://www.emfink.net/EmploymentLaw/assets/materials/Confidentiality_Competition/Krueger-Posner.pdf.

Kwoka, John. 2014. *Mergers, Merger Control, and Remedies: A Retrospective Analysis of U.S. Policy*. MIT Press. Google-Books-ID: qcYQBgAAQBAJ.

Lamadon, Thibaut, Magne Mogstad, and Bradley Setzler. 2019. "Imperfect Competition, Compensating Differentials and Rent Sharing in the U.S. Labor Market." Working Paper 25954, National Bureau of Economic Research. URL https://www.nber.org/papers/w25954. Series: Working Paper Series.

———. 2022. "Imperfect Competition, Compensating Differentials, and Rent Sharing in the US Labor Market." *American Economic Review* 112 (1):169–212. URL https://pubs.aeaweb.org/doi/10.1257/aer.20190790.

Lavetti, Kurt. 2023. "Compensating Wage Differentials in Labor Markets: Empirical Challenges and Applications." *Journal of Economic Perspectives* 37 (3):189–212. URL https://www.aeaweb.org/articles?id=10.1257/jep.37.3.189.

Lipsitz, Michael and Mark J Tremblay. 2021. "Noncompete agreements and the welfare of consumers." *Available at SSRN 3975864* .

Loecker, Jan De and Frederic Warzynski. 2012. "Markups and firm-level export status." *American economic review* 102 (6):2437–2471.

Manning, Alan. 2011. "Chapter 11 - Imperfect Competition in the Labor Market." In *Handbook of Labor Economics*, vol. Volume 4, Part B. Elsevier, 973–1041. URL http://www.sciencedirect.com/science/article/pii/S0169721811024099.

———. 2013. *Monopsony in motion: Imperfect competition in labor markets*. Princeton University Press.

———. 2021. "Monopsony in Labor Markets: A Review." *ILR Review* 74 (1):3–26. URL https://doi.org/10.1177/0019793920922499. Publisher: SAGE Publications Inc.

Marinescu, Ioana and Herbert Hovenkamp. 2019. "Anticompetitive Mergers in Labor Markets." *Indiana Law Journal* 94 (3):1031. URL https://www.repository.law.indiana.edu/ilj/vol94/iss3/5.

Marinescu, Ioana, Ivan Ouss, and Louis-Daniel Pape. 2021. "Wages, hires, and labor market concentration." *Journal of Economic Behavior & Organization* 184:506–605. URL https://www.sciencedirect.com/science/article/pii/S0167268121000548.

Marinescu, Ioana and Roland Rathelot. 2018. "Mismatch Unemployment and the Geography of Job Search." *American Economic Journal: Macroeconomics* 10 (3):42–70. URL https://www.aeaweb.org/articles?id=10.1257/mac.20160312.

Marinescu, Ioana Elena and Eric A. Posner. 2019. "Why Has Antitrust Law Failed Workers?" *SSRN Electronic Journal* URL https://www.ssrn.com/abstract=3335174.

Marx, Matt. 2018. "Reforming Non-Competes to Support Workers." Policy Proposal 2018-04, The Hamilton Project. URL https://www.hamiltonproject.org/assets/files/reforming_noncompetes_support_workers_marx_policy_proposal.pdf.

Matvos, Gregor and Michael Ostrovsky. 2008. "Cross-ownership, returns, and voting in mergers." *Journal of Financial Economics* 89 (3):391–403.

Meccheri, Nicola. 2009. "A note on noncompetes, bargaining and training by firms." *Economics Letters* 102 (3):198–200. URL https://www.sciencedirect.com/science/article/pii/S0165176509000172.

Mertens, Matthias. 2022. "Micro-mechanisms behind declining labor shares: Rising market power and changing modes of production." *International Journal of Industrial Organization* 81:102808.

Miller, Nathan, Steven Berry, Fiona Scott Morton, Jonathan Baker, Timothy Bresnahan, Martin Gaynor, Richard Gilbert, George Hay, Ginger Jin, Bruce Kobayashi, Francine Lafontaine, James Levinsohn, Leslie Marx, John Mayo, Aviv Nevo, Ariel Pakes, Nancy Rose, Daniel Rubinfeld, Steven Salop, Marius Schwartz, Katja Seim, Carl Shapiro, Howard Shelanski, David Sibley, Andrew Sweeting, and Marta Wosinska. 2022. "On the misuse of regressions of price on the HHI in merger review." *Journal of Antitrust Enforcement* 10 (2):248–259. URL https://doi.org/10.1093/jaenfo/jnac009.

Moen, Espen R. 1997. "Competitive Search Equilibrium." *Journal of Political Economy* 105 (2):385–411. URL https://www.journals.uchicago.edu/doi/abs/10.1086/262077. Publisher: The University of Chicago Press.

O'brien, Daniel P and Steven C Salop. 1999. "Competitive effects of partial ownership: Financial interest and corporate control." *Antitrust LJ* 67:559.

OIRA, OMB. 2023. "Guidance on Accounting for Competition Effects When Developing and Analyzing Regulatory Actions." .

Okudaira, Hiroko, Miho Takizawa, and Kenta Yamanouchi. 2019. "Minimum wage effects across heterogeneous markets." *Labour Economics* 59:110–122.

Posner, Eric and Eric Weyl. 2018. *Radical markets: Uprooting capitalism and democracy for a just society*. Princeton University Press.

Posner, Eric A, Fiona M Scott Morgan, and E Glen Weyl. 2016. "A proposal to limit the anti-competitive power of institutional investors." *Antitrust LJ* 81:669.

Posner, Eric A., George G. Triantis, and Alexander J. Triantis. 2004. "Investing in Human Capital: The Efficiency of Covenants Not to Compete." *University of Chicago Law & Economics, Olin Working Paper* (No. 137):1–35. URL https://papers.ssrn.com/abstract=285805.

Postel-Vinay, Fabien and Jean-Marc Robin. 2002. "The Distribution of Earnings in an Equilibrium Search Model with State-Dependent Offers and Counteroffers*." *International Economic Review* 43 (4):989–1016. URL http://onlinelibrary.wiley.com/doi/10.1111/1468-2354.t01-1-00045/abstract.

Prager, Elena and Matt Schmitt. 2021. "Employer Consolidation and Wages: Evidence from Hospitals." *American Economic Review* 111 (2):397–427. URL https://www.aeaweb.org/articles?id=10.1257/aer.20190690.

Robinson, Joan. [1933] 1966. "Imperfect Competition." In *An Essay on Marxian Economics*, edited by Joan Robinson. Palgrave Macmillan UK, 73–81. URL https://doi.org/10.1007/978-1-349-15228-5_9.

Rosen, Sherwin. 1974. "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition." *Journal of Political Economy* 82 (1):34–55. URL https://www.journals.uchicago.edu/doi/10.1086/260169.

———. 1986. "Chapter 12 The theory of equalizing differences." In *Handbook of Labor Economics*, vol. 1. Elsevier, 641–692. URL https://www.sciencedirect.com/science/article/pii/S1573446386010155.

Roussille, Nina and Benjamin Scuderi. 2023. "Bidding for Talent: Equilibrium Wage Dispersion on a High-Wage Online Job Board."

Rubin, Paul H. and Peter Shedd. 1981. "Human Capital and Covenants Not to Compete." *The Journal of Legal Studies* 10 (1):93–110. URL https://www.journals.uchicago.edu/doi/10.1086/467672.

Rudanko, Leena. 2023. "Firm Wages in a Frictional Labor Market." *American Economic Journal: Macroeconomics* 15 (1):517–550. URL https://pubs.aeaweb.org/doi/10.1257/mac.20200440.

Shi, Liyan. 2023. "Optimal Regulation of Noncompete Contracts." *Econometrica* 91 (2):425–463. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA18128. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA18128.

Sockin, Jason. 2021. "Show Me the Amenity: Are Higher-Paying Firms Better All Around?" URL https://papers.ssrn.com/abstract=3957002.

Sokolova, Anna and Todd Sorensen. 2021. "Monopsony in Labor Markets: A Meta-Analysis." *ILR Review* 74 (1):27–55. URL https://doi.org/10.1177/0019793920965562.

Starr, Evan. 2019. "Consider This: Training, Wages, and the Enforceability of Covenants Not to Compete." *ILR Review* 72 (4):783–817. URL http://journals.sagepub.com/doi/10.1177/0019793919826060.

Starr, Evan P., J.J. Prescott, and Norman D. Bishara. 2021. "Noncompete Agreements in the US Labor Force." *The Journal of Law and Economics* 64 (1):53–84. URL https://www.journals.uchicago.edu/doi/10.1086/712206.

U.S. Department of Justice and the Federal Trade Commission. 2023. "Draft FTC-DOJ Merger Guidelines for Public Comment." Draft - For Public Comment Purposes FTC-2023-0043-0001. URL https://www.ftc.gov/system/files/ftc_gov/pdf/p859910draftmergerguidelines2023.pdf.

Verboven, Frank. 1996. "The nested logit model and representative consumer theory." *Economics Letters* 50 (1):57–63.

Wu, Tim. 2018. *The Curse of Bigness: Antitrust in the New Gilded Age*. Columbia Global Reports. URL https://www.jstor.org/stable/j.ctv1fx4h9c.

Yeh, Chen, Claudia Macaluso, and Brad Hershbein. 2022. "Monopsony in the US Labor Market." *American Economic Review* 112 (7):2099–2138. URL https://www.aeaweb.org/articles?id=10.1257/aer.20200025.

Zingales, Luigi. 2017. "Towards a Political Theory of the Firm." *The Journal of Economic Perspectives* 31 (3):113–130. URL https://www.jstor.org/stable/44321282. Publisher: American Economic Association.