

Firm Market Power and the Earnings Distribution

Douglas A. Webber ^{*†}

11-22-11

Abstract

Using the Longitudinal Employer Household Dynamics (LEHD) data from the United States Census Bureau, I compute firm-level measures of labor market (monopsony) power. To generate these measures, I extend the dynamic model proposed by Manning (2003) and estimate the labor supply elasticity facing each private non-farm firm in the US. While a link between monopsony power and earnings has traditionally been assumed, I provide the first direct evidence of the positive relationship between a firm's labor supply elasticity and the earnings of its workers. I also contrast the semistructural method with the more traditional use of concentration ratios to measure a firm's labor market power. In addition, I provide several alternative measures of labor market power which account for potential threats to identification such as endogenous mobility. Finally, I construct a counterfactual earnings distribution which allows the effects of firm market power to vary across the earnings distribution.

I estimate the average firm's labor supply elasticity to be 1.08, however my findings suggest there to be significant variability in the distribution of firm market power across US firms, and that dynamic monopsony models are superior to the use of concentration ratios in evaluating a firm's labor market power. I find that a one-unit increase in the labor supply elasticity to the firm is associated with wage gains of between 5 and 18 percent. While nontrivial, these estimates imply that firms do not fully exercise their labor market power over their workers. Furthermore, I find that the negative earnings impact of a firm's market power is strongest in the lower half of the earnings distribution, and that a one standard deviation increase in firms' labor supply elasticities reduces the variance of the earnings distribution by 9 percent.

*This research uses data from the Census Bureau's Longitudinal Employer Household Dynamics Program, which was partially supported by the National Science Foundation Grants SES-9978093, SES-0339191 and ITR-0427889; National Institute on Aging Grant AG018854; and grants from the Alfred P. Sloan Foundation. Any opinions and conclusions expressed herein are those of the author and do not necessarily represent the views of the U.S. Census Bureau, its program sponsors or data providers, or of Cornell University. All results have been reviewed to ensure that no confidential information is disclosed.

†I have greatly benefited from the advice of John Abowd, Francine Blau, Ron Ehrenberg, Kevin Hallock, George Jakubson, and Alan Manning. I would also like to sincerely thank Henry Hyatt, J. Catherine Maclean, Matt Masten, Erika Mcentarfer, Ben Ost, and Michael Strain for their many helpful comments.

1 Introduction

There is good reason to believe that some firms have non-trivial power in the labor market, that not all firms act as price takers and pay their employees the prevailing market wage. Intuitively, most would not switch jobs following a wage cut of one cent, and we would not expect a firm which raises wages by a small amount to suddenly have an infinite stream of workers. So it becomes an empirical question of whether the departure from perfect competition is meaningful; whether perfect competition is a good approximation for our economy, or whether a model with substantial frictions fits better.

The existence of significant firm effects in wage regressions, even after controlling for detailed person and industry characteristics, is cited as strong suggestive evidence of firm market power (Abowd et al., 1999; Goux and Maurin, 1999). For instance, Goux and Maurin (1999) conclude that on average firm effects alter an individual's wage by more than 20 percent. Furthermore, they find these firm effects are related more to firm characteristics such as size rather than productivity, implying that the firm effects are not simply absorbing workers' unmeasured marginal product of labor.

Estimating the degree of wage competition in the labor market is important for both theoretical research and policy analysis. Since perfect competition is a standard feature in many models of the labor market, evidence of significant distortions in the labor market would suggest labor economists should reevaluate the perfect competition assumption and its implications in their models. From a policy perspective, the degree of imperfect competition can drastically change the effects of institutions such as the minimum wage (Card and Krueger, 1995) or unions (Feldman and Scheffler, 1982).

While the industrial organization literature has theoretically and empirically modeled similar frictions in the product market, there has been comparatively less work done to account for distortions of the labor market. This is primarily due to the comparative lack of rich labor market data (such as linked employer-employee data) versus product market data. Most of the theoretical work done on this topic resides in the search theory literature,

with major contributions coming from Burdett and Mortensen (1998) and Shimer (2005) to name a few¹. This line of research has given rise to a "new monopsony" literature, popularized by Alan Manning's (Manning, 2003) careful analysis of labor-related topics absent the assumption of perfect competition. The new monopsony model of the labor market views a firm's market power as derived from search frictions rather than solely geographic power as in a classic monopsony model. These search frictions originate from imperfections in the labor market such as imperfect information about available jobs, worker immobility, or heterogeneous preferences.

Even if the existence of monopsony power is accepted, estimating the degree of market power possessed by a firm is not a simple task. Economists since Bunting (1962) have searched for empirical evidence of monopsony, with the predominant method being the use of concentration ratios, the share of a labor market which a given firm employs. The most commonly examined market in the empirical monopsony literature has been that of nurses in hospitals (Hurd, 1973; Feldman and Scheffler, 1982; Hirsch and Schumacher, 1995; Link and Landon, 1975; Adamache and Sloan, 1982; Link and Settle, 1979). This market lends itself to monopsony because nurses have a highly specific form of human capital and there are many rural labor markets where hospitals are the dominant employer. Despite the relatively large literature on this narrow labor market, the concentration ratio approach has yielded mixed results and no clear consensus.

More recently, studies have attempted to directly estimate the average slope of the labor supply curve faced by the firm, which is a distinct concept from the market labor supply elasticity². Studying the market for nurses, Sullivan (1989) finds evidence of monopsony using a structural approach to measure the difference between nurses' marginal product of labor and their wages. Examining another market commonly thought to be monopsonistic, the market for schoolteachers, Ransom and Sims (2010) instrument wages with collectively

¹See Mortensen (2003) or Rogerson et al. (2005) for a review of this literature

²The market labor supply elasticity corresponds to the decision of a worker to enter the labor force, while the labor supply elasticity to the firm corresponds to the decision of whether to supply labor to a particular firm. This paper focuses on the firm-level decision.

bargained pay scales and estimate a labor supply elasticity between 3 and 4. In a novel approach using German administrative data, Schmieder (2010) finds evidence of a positive sloping labor supply curve through an analysis of new establishments.

Using a dynamic approach similar to this study, Ransom and Oaxaca (2010) and Hirsch et al. (2010) both separately estimate the labor supply elasticities to the firm of men and women, each finding strong evidence of monopsonistic competition. Ransom and Oaxaca (2010) use data from a chain of grocery stores, and find labor supply elasticities of about 2.5 for men and 1.6 for women. Hirsch et al. (2010) uses administrative data from Germany to estimate elasticities ranging from 2.5-3.6 and 1.9-2.5 for men and women respectively. Applying this approach to survey data, Manning (2003) finds labor supply elasticities ranging from 0.68 in the NLSY to 1.38 in the PSID. In a developing country context, Brummund (2011) finds strong evidence of monopsony in Indonesian labor markets, estimating labor supply elasticities between .6 and 1.

Utilizing data from the US Census Bureau's Longitudinal Employer Household Dynamics (LEHD) program, I estimate the market-level average labor supply elasticity faced by firms in the US economy, similar to the Hirsch et al. (2010) study using German data. I then extend the approach to estimate firm-level labor supply elasticities. This is accomplished through an extension to the dynamic model of labor supply proposed by Manning (2003). This method allows me to examine the effects of monopsonistic competition on the earnings distribution in great detail, and contributes to the existing literature in a number of ways. First, it is the first examination of monopsony power using comprehensive administrative data from the US. Second, my particular empirical strategy allows me to examine the distribution of monopsony power which exists in the US, and to provide the first direct evidence on the negative impact of a firm's market power on earnings. I compare the performance of the market power measures derived in this study to that of the more traditional concentration ratio to illustrate the significant contribution of the new monopsony models. Finally, I construct a counterfactual earnings distribution in which firms' market power is reduced

in order to demonstrate the impact of imperfect competition on the shape of the earnings distribution.

I estimate the average labor supply elasticity to the firm to be approximately 1.08. Estimates in this range are robust to various modeling assumptions and corrections for endogenous mobility. Furthermore, I find evidence of substantial heterogeneity in the market power possessed by firms, ranging from negligible to highly monopsonistic. While a link between monopsony power and wages has traditionally been assumed (Pigou, 1924), I provide the first direct evidence of a positive relationship between a firm’s labor supply elasticity and the earnings of its workers, estimating that a one-unit increase is associated with a decrease of between .05 and .18 log earnings. I demonstrate that the effect of monopsony power is not constant across workers: unconditional quantile regressions imply that impacts are largest among low paid and negligible among high paid workers. Finally, implications in the inequality literature are addressed through the construction of a counterfactual earnings distribution, which implies that a one standard deviation increase of each firm’s labor supply elasticity would decrease the variance of earnings distribution by 9 percent.

The paper is organized as follows, Section 2 describes the definition of market power utilized in this study. Section 3 lays out the theoretical foundation for this study. The data and methods are described in Section 4. Section 5 presents the results and sensitivity analyses, and Section 6 concludes.

2 Discussion of Monopsony Power

The concept of “monopsony” was first defined and explored as a model by Robinson (1933). In her seminal work, Robinson formulated the analysis which is still taught in undergraduate labor economics courses. Monopsony literally means “one buyer”, and although the term is most often used in a labor market context, it can also refer to a firm which is the only buyer of an input.

It should be pointed out that in the “new monopsony” framework, the word monopsony is synonymous with the following phrases: monopsonistic competition, imperfect competition, finite labor supply elasticity, or upward sloping labor supply curve to the firm. While the classic monopsony model is based on the idea of a single firm as the only outlet for which workers can supply labor, the new framework defines monopsony as any departure from the assumptions of perfect competition. Additionally, the degree of monopsonistic competition may vary significantly across labor markets, and even across firms within a given labor market.

In order to think about what determines a firm’s monopsony power, we must consider why we do not observe the predicted behavior from a perfectly competitive model. What gives a firm flexibility in offering a wage rather than being forced to offer the market wage? Put another way, why do we not observe workers jumping from job to job whenever they observe a higher paying opportunity for which they are qualified?

One of the most prominent reasons is that the typical worker does not have a continuous stream of job offers (this point will be discussed further in the theoretical model section). This source of monopsony power has roots in the classic monopsony framework in that, all else held constant, workers in labor markets with more firms are likely to have a greater number of offers. However, this idea takes an overly simplistic view of the boundaries of a given labor market. Most employers are likely operating in many labor markets at any given time. A prestigious university may be competing in a national or international labor market for professors, a regional labor market for its high-level administrators and technical staff, and a local labor market for the low-level service workers. Even if the arrival rate of job offers were the only source of monopsony power, it seems that geographic modeling alone would do a poor job of measuring that power. Another source of monopsony power is imperfect information about job openings (McCall, 1970; Stigler, 1962), which is not completely distinct from the arrival rate of job offers since a decrease in information can cause a reduction in job offers. This is a particularly compelling example since studies

such as Hoffer and Murphy (1992) and Polachek and Robst (1998) estimate that imperfect information about job prospects depresses wages by approximately ten percent.

The costs (both monetary and psychic) associated with changing jobs can also be thought of as giving market power to the firm. Moving costs are typically thought of as a short run cost, particularly when a worker is young. However these costs can grow significantly when a worker has a family and roots in a community. Consider the scenario of a dual-career family. Two job offers will be needed to induce either of the partners to move, a fact which gives significant bargaining power to the employers of each partner, particularly the one who is paid less. Additionally, changing jobs means that workers must adjust to a new system which will require at least a small degree of learning on the job.

Firm specific human capital also can be thought of as giving market power to the firm, since there is in effect a barrier to leaving a firm when an individual's firm specific capital is large relative to their general human capital. In fact, Wasmer (2006) concludes that markets with substantial search frictions induce workers to overinvest in firm specific human capital.

Reputation costs likely also play a large role in the mobility of workers. Potential employers would be very suspicious of hiring a worker who changes jobs the moment he is offered any wage increase. For all of these reasons, and likely many more, workers must be selective with the wage offers they choose to accept, thus leading to a labor market with substantial frictions.

As discussed in Manning (2010), another way to think about imperfect competition in the labor market is in terms of the rents received by the employee and the employer. On the worker's side, the rents to a given job match would be the difference between the current wage (utility) and the worker's opportunity cost, either a wage (utility) from a different firm or unemployment benefits. Studies such as Jacobson et al. (1993) implicitly estimate these rents by exploring the impacts of exogenous job destruction. This literature estimates wage losses of 20-30 percent, implying significant rents to employees from a given job match. From the employer's perspective, the rents from the i th job match are the difference between

$(MP_i - w_i)$ and $(MP_j - w_j)$, where j is the next worker who would be hired if worker i leaves the firm. This is a harder quantity to measure empirically, but can be approximated (assuming that the marginal product is the same for workers i and j) by hiring and training costs. The estimates of hiring and training costs as a fraction of total wages paid tend to be in the range of 3-7 percent (Oi, 1962; Abowd and Kramarz, 2003). The ratio of worker rents to employer rents can be thought of as a measure of the firm's market power. If the worker's opportunity cost is high relative to her employer's opportunity cost, then the employer will be able to extract a large amount of the surplus from the job match. However, if the converse is true, the worker will be in the position of power.

A relatively new branch of labor economics which focuses on the initial labor market conditions when a worker enters the labor market may also provide insight into the mobility of workers. A number of studies (Oyer, 2006, 2008; Genda and Kondo, 2010; Kahn, 2010) find persistent and negative wage effects from entering the labor market in a bad economy, lasting for at least 20 years. These persistent effects provide further evidence that there are significant long-run frictions in the economy.

Finally, while a worker's earnings represent an important market outcome, it is important to remember that wages make up only a part of the total "compensation" to the worker. The true quality of a job match has many dimensions, such as benefits, working conditions, and countless other compensating differentials. The interaction of monopsony with these non-wage goods should be explored in future research.

3 Theoretical Model

A central feature of perfect competition is the law of one wage, that all workers of equal ability should be paid the same market clearing wage. In an attempt to explain how wage dispersion can indeed be an equilibrium outcome, Burdett and Mortensen (1998) develop a model of the economy in which employers post wages based on the wage-posting behavior

of competing employers. Even assuming equal ability for all workers, wage dispersion is an equilibrium outcome as long as one assumes that the arrival rate of job offers is positive but finite (perfect competition characterizes the limiting case, as the arrival rate tends to infinity). While I do not explicitly estimate the Burdett and Mortensen model in this paper, the intuition of monopsony power derived from search frictions is central to this study. See Kuhn (2004) for a critique of the use of equilibrium search models in a monopsony context.

The Burdett and Mortensen model of equilibrium wage dispersion

Assume there are M_t equally productive workers (where productivity is given by p), each gaining utility b from leisure. Further assume there are M_e constant returns to scale firms which are infinitesimally small when compared to the entire economy. A firm sets wage w to maximize steady-state profits $\pi = (p-w)N(w)$ where $N(w)$ represents the supply of labor to the firm. Also define $F(w)$ as the cdf of wage offers observed in the economy, and $f(w)$ is the corresponding pdf. All workers within a firm must be paid the same wage. Employed workers will accept a wage offer w' if it is greater than their current wage w , and non-employed workers will accept w' if $w' \geq b$ where b is their reservation wage. Wage offers are drawn randomly from the distribution $F(w)$, and arrive to all workers at rate λ . Assume an exogenous job destruction rate δ , and that all workers leave the job market at rate δ to be replaced in nonemployment by an equivalent number of workers. R^N denotes The recruitment flow and separation rate functions are given by:

$$R(w) = R^N + \lambda \int_0^w f(x)N(x)dx \tag{1}$$

$$s(w) = \delta + \lambda(1 - F(w)) \tag{2}$$

Burdett and Mortensen (1998), or alternatively Manning (2003), show that in this economy, as long as λ is positive and finite, there will be a nondegenerate distribution of wages even when all workers are equally productive. As λ tends to zero, the wage distribution will

collapse to the monopsony wage, which in this particular economy would be the reservation wage b . As λ tends to infinity the wage distribution will collapse to the perfectly competitive wage, the marginal product of labor p .

Note that the following primarily relies on the model presented in Manning (2003), and incorporates a key insight from the recent working paper by Depew and Sorensen (2011) to derive the least restrictive formula for the labor supply elasticity facing the firm currently in the literature. We can recursively formulate the supply of labor to a firm with the following equation, where $R(w)$ is the flow of recruits to a firm and $s(w)$ is the separation rate.

$$N_t(w) = N_{t-1}(w)[1 - s_{t-1}(w)] + R_{t-1}(w) \quad (3)$$

Equation (3) formalizes the definitionally true statement that a firm's employment this period is equal to the fraction of workers from last period who stay with the firm plus the number of new recruits. Noting that $N_t = \gamma N_{t-1}$ where γ is the rate of employment growth between period $t-1$ and t , we can rewrite Equation (3) as

$$N_t(w) = \frac{R_t(w)}{1 - (1 - s_t(w))^{\frac{1}{\gamma_t}}} \quad (4)$$

Taking the natural log of each side, multiplying by w , and differentiating we can write the elasticity of labor supply, ε , at time t as a function of the long-run elasticities of recruitment and separations, as well as the contemporary separation and growth rates.

$$\varepsilon_t = \varepsilon_R - \varepsilon_S \frac{s_t(w)}{\gamma_t + s_t(w) - 1} \quad (5)$$

We can further decompose the recruitment and separation elasticities in the following way

$$\varepsilon_t = \theta^R \varepsilon_R^E + (1 - \theta^R) \varepsilon_R^N - \theta^S \varepsilon_S^E \frac{s_t^E(w)}{\gamma_t + s_t^E(w) - 1} - (1 - \theta^S) \varepsilon_S^N \frac{s_t^N(w)}{\gamma_t + s_t^N(w) - 1} \quad (6)$$

Where the elasticity of recruitment has been broken down into the elasticity of recruitment of workers from employment (ε_R^E) and the elasticity of recruitment of workers from nonemployment (ε_R^N). Similarly the elasticity of separation has been decomposed into the elasticity of separation to employment (ε_S^E) and the elasticity of separation to nonemployment (ε_S^N). θ^R and θ^S represent the share of recruits from employment and the share of separations to employment respectively.

While there are established methods for estimating separation elasticities with standard job-flow data, recruitment elasticities are not identified without detailed information about every job *offer* a worker receives. Therefore, it would be helpful to express the elasticities of recruitment from employment and noemployment as functions of estimable quantities.

Looking first at the elasticity of recruitment from employment, we can write the recruitment from employment function and its derivative as

$$R^E(w) = \lambda \int_0^w f(x)N(x)dx \quad (7)$$

$$\frac{\partial R^E(w)}{\partial w} = \lambda f(w)N(w) \quad (8)$$

Combining Equations (4), (7), and (8), along with the definition of an elasticity ($\varepsilon_R^E = \frac{w}{R^E(w)} \frac{\partial R^E(w)}{\partial w}$), we get:

$$\varepsilon_R^E = \frac{w\lambda f(w)}{1 + \frac{s^E(w)}{\gamma_t} - \frac{1}{\gamma_t}} \quad (9)$$

In dealing with the numerator, note that the the derivative of the separation to employment function, $s^E(w) = \lambda(1 - F(w))$, is

$$\frac{\partial s^E(w)}{\partial w} = -\lambda f(w) \quad (10)$$

Combining equations (9), (10), and the definition of an elasticity ($\varepsilon_s^E = \frac{w}{s^E(w)} \frac{\partial s^E(w)}{\partial w}$), we

can write the elasticity of recruitment from employment as a function of estimable quantities:

$$\varepsilon_R^E = \frac{-\varepsilon_R^E s_t^E(w)}{1 + \frac{s_t^E(w)}{\gamma_t} - \frac{1}{\gamma_t}} \quad (11)$$

Next, Manning (2003, p. 100) notes that the elasticity of recruitment from nonemployment can be written as

$$\varepsilon_R^N = \varepsilon_R^E - w\theta^R(w)/\theta^R(w)(1 - \theta^R(w)) \quad (12)$$

This is derived from the simple definition of θ^R , the share of total recruits which come from employment, which implies $R^N = R^E(1 - \theta^R)/\theta^R$, where R^N and R^E are the recruits from nonemployment and employment respectively. Taking the natural log of each side of this relation and differentiating yields the relation depicted in Equation (12). The second term on the right-hand side of Equation (12) can be thought of as the bargaining premium that an employee receives from searching while currently employed. Thus, the labor supply elasticity to the firm can be written as a function of both separation elasticities, the premium to searching while employed, and the calculated separation and growth rates. To my knowledge, no other study has estimated this model before.

In an economy where the arrival rate of job offers is finite (and thus the labor supply elasticity is finite) firms are not bound by market forces to pay workers their marginal product of labor. The model presented above implies that, even in a world where all firms and individuals are identical, a decrease in the arrival rate of job offers will both lower the average wage and increase inequality. To see how a firm's labor supply elasticity affects the wage it pays, consider a profit-maximizing firm which faces the following objective function:

$$\underset{w}{Max} \Pi = pQ(L) - wL(w) \quad (13)$$

P is the price of the output produced according to the production function Q. The

choice of wage w determines the labor supplied to the firm L . Taking first order conditions, substituting $\varepsilon = \frac{w}{L(w)} \frac{\partial L(w)}{\partial w}$, and solving for w yields:

$$w = \frac{pQ'(L)}{1 + \frac{1}{\varepsilon}} \quad (14)$$

The numerator in Equation (14) is simply the marginal product of labor, and ε is the labor supply elasticity faced by the firm. It is easy to see that in the case of perfect competition ($\varepsilon = \infty$) that the wage is equal to the marginal product of labor, but the wage is less than then marginal product for all $0 < \varepsilon < \infty$.

Every empirical study in the new monopsony literature attempts to estimate the labor supply elasticity to the firm at the market level. In other words, they measure the (firm-size weighted) average slope of each firm's supply curve in the market. In a highly competitive market we would expect these elasticities to be very large numbers. Among the contributions of this paper is to separately estimate each firm's labor supply elasticity rather than a market average.

4 Data and Methodology

Data

The Longitudinal Employer Household Dynamics (LEHD) data are built primarily from Unemployment Insurance (UI) wage records, which cover approximately 98 percent of wage and salary payments in private sector non-farm jobs. Information about the firms is constructed from the Quarterly Census of Employment and Wages (QCEW). The LEHD infrastructure allows users to follow both workers and firms over time, as well as to identify workers who share a common employer. Firms in these data are defined at the state level, which means that a Walmart in Florida and a Walmart in Georgia would be considered to be different firms. However, all Walmarts in Florida are considered to be part of the same firm. These

data also include demographic characteristics of the worker and basic firm characteristics, obtained through administrative record and statistical links. For a complete description of these data, see Abowd et al. (2009).

My sample consists of quarterly observations on earnings and employment for 47 states between 1985 and 2008³. I make several sample restrictions in an attempt to obtain the most economically meaningful results, analyses without these restrictions are presented as a robustness check later in the paper. These restrictions are necessary in large part because the earnings data are derived from tax records, and thus any payment made to an individual, no matter how small, will appear in the sample. As a consequence, there are many “job spells” which appear to last only one quarter, but are in fact one-time payments which do not conform with the general view of a job match between a firm and worker.

First, I only include an employment spell in the sample if at some point it could be considered the dominant job, defined as paying the highest wage of an individual’s jobs in a given quarter⁴. I also remove all spells which span fewer than three quarters.⁵ This sample restriction is related to the construction of the earnings variable. Since the data do not contain information on when in the quarter an individual was hired/separated, the entries for the first and last quarters of any employment spell will almost certainly underestimate the quarterly earnings rate (unless the individual was hired on the first day or left employment on the last day of a quarter). Thus, in order to get an accurate measurement of the earnings rate I must observe an individual in at least one quarter other than the first or last of an employment spell. I remove job spells which have average earnings greater than \$1 million per quarter and less than \$100 per quarter, which corresponds approximately to the top and

³The states not in the sample are Connecticut, Massachusetts, and New Hampshire. Not all states are in the LEHD infrastructure for the entire time-frame, but once a state enters it is in the sample for all subsequent periods. Figure 1 presents the coverage level of the US economy reproduced from Abowd and Vilhuber (2011).

⁴This formulation allows an individual to have more than one dominant job in a given quarter. The rationale behind this definition is that I wish to include all job spells where the wage is important to the worker. The vast majority of job spells in my sample, 89.9 percent, have 0 or 1 quarters of overlap with other job spells. Restricting the dominant job definition to only allow one dominant job at a given time does not alter the reported results.

⁵The relaxation of this assumption does not appreciably alter any of the reported results.

bottom 1 percent of observations

Additionally, I limit the analysis to firms with 100 total employment spells of any length over the lifespan of the firm. For the full-economy monopsony model, these sample restrictions yield a final sample of approximately 149,710,000 unique individuals who had 325,630,000 total employment spells at 670,000 different firms. Additionally, for analyses using the firm-level measure of the labor supply elasticity, only firms which have greater than 25 separations to employment, 25 separations to unemployment, and 25 recruits from employment over the lifespan of the firm are considered. This reduces the analysis sample to approximately 121,190,000 unique individuals having 267,310,000 employment spells at 340,000 unique firms.

Empirical Strategy

The primary reason for the small empirical literature on monopsony is a lack of high quality data. In order to identify a firm's market power, the researcher must have a credible firm-level instrument for each firm studied or detailed employer-employee linked data to identify worker flows. I employ the latter approach in this study since finding a credible instrument for nearly every firm in the US is unlikely. The construction of the market power measures most closely represents an augmented firm-level implementation of the methodology proposed in Manning (2003).

I first describe in detail how the market power measures are calculated, followed by a description of how they are used to examine the US earnings distribution.

Location-Based Measures

I construct an overall measure of the percent of the industry-specific labor market that each firm employs (Number of workers at firm i /number of workers in firm i 's county and in firm i 's industry) using North American Industry Classification System (NAICS) industry definitions. While this variable is far from a perfect measure of an employer's power to

set wages, it has several advantages over the dynamic measures to be used later in the paper. Both the construction of these measures and the regression estimates using them are transparent. Endogeneity, misspecified equations, etc. are of less concern in the construction of these labor concentration measures, and the interpretation of the regression coefficients on these variables is straightforward. This analysis corresponds to the traditional concentration ratio approach of analyzing labor market power.

Dynamic Measure

The simplest way to estimate the labor supply elasticity to the firm would be to regress the natural log of firm size on the natural log of firm wages. However, even when controlling for various demographic characteristics, this is deemed to produce a potentially biased estimate⁶. I therefore rely on estimating parameters presented in the theoretical section which are plausibly identified, and then combine them using results from Manning (2003) and equation (6) to produce an estimate of the labor supply elasticity to the firm.

To my knowledge, only Hirsch et al. (2010) has used a similar, but considerably more restrictive, method with administrative data which yielded an economy-wide estimate of the average labor supply curve facing the firm. Manning (2003) also estimates an economy-wide measure of the degree of monopsony using surveys such as the National Longitudinal Survey of Youth (NLSY) 1979. One of the major contributions of this paper is that I estimate the labor supply elasticities for each firm, rather than the average over the whole economy. Additionally, these prior studies imposed a steady-state assumption on their model, which the model in this paper does not impose. Estimating the labor supply elasticities at the firm level does have several advantages. First, the estimation of each of the elasticity components is much more flexible than even the least constrained specifications of Hirsch et al. (2010). Second, I will be able to use the measures as an explanatory variable, and can test a number

⁶The firm size-wage premium is a well known result in the labor economics literature, and is often attributed to non-monopsony related factors such as economies of scale increasing the productivity, and thus the marginal product, of workers at large firms

of different models. Finally, I will be able to examine the effect of market power on earnings at each point in the market power distribution, rather than examining only the average effect. This is particularly important because theory predicts significant nonlinear effects relating to the labor supply elasticity and a firm's ability to mark down wages (Pigou, 1924). However, this strategy has the drawback that I am unable to estimate the relevant parameters, and thus the labor supply elasticity, for the smallest firms (sample restrictions are discussed in the data section).

According to the results presented in the theoretical model section, three quantities must be estimated in order to construct the labor supply elasticity measure, $(\varepsilon_S^E, \varepsilon_S^N$ and $w\theta^{R'}(w)/\theta^R(w)(1 - \theta^R(w))$), as well as the calculated separation and growth rates for each firm. Each of the following models will be run separately for every firm in the sample (as well as on the whole sample for comparison purposes), where the unit of observation is an employment spell, thus one individual can appear in multiple firm's models. Looking first at the separation elasticities, I model separations to nonemployment as a Cox proportional hazard model given by

$$\lambda^N(t|\beta^{N,sep}\log(earnings)_i + X_i\gamma^{N,sep}) = \lambda_0(t) \exp(\beta^{N,sep}\log(earnings)_i + X_i\gamma^{N,sep}) \quad (15)$$

where $\lambda()$ is the hazard function, λ_0 is the baseline hazard, t is the length of employment, $\log(earnings)$ is the natural log of individual i 's average quarterly earnings, and X is a vector of explanatory variables including gender, race, age, education, and year control variables (industry controls are also included in the full-economy model). While the entire sample will be used, workers who transition to a new employer or who are with the same employer at the end of the data series are considered to have a censored employment spell. In this model, the parameter β represents an estimate of the separation elasticity to nonemployment. In an analogous setting, I model separations to employment as

$$\lambda^E(t|\beta^{E,sep}\log(earnings)_i + X_i\gamma^{E,sep}) = \lambda_0(t) \exp(\beta^{E,sep}\log(earnings)_i + X_i\gamma^{E,sep}) \quad (16)$$

with the only difference being that the sample is restricted to those workers who do not have a job transition to nonemployment. As before, β represents an estimate of the separation elasticity to employment. To estimate the third quantity needed for equation (6), $w\theta^R(w)/\theta^R(w)(1 - \theta^R(w))$, Manning (2003) shows that this is equivalent to the coefficient on log earnings when estimating the following logistic regression

$$P_{rec} = \frac{\exp(\beta^{E,rec}\log(earnings)_i + X_i\gamma^{E,rec})}{1 + \exp(\beta^{E,rec}\log(earnings)_i + X_i\gamma^{E,rec})} \quad (17)$$

where the dependent variable takes a value of 1 if a worker was recruited from employment and 0 if they were recruited from nonemployment. To enable this coefficient to vary over time, log earnings is interacted with time dummies. The same explanatory variables used in the separation equations are used in this logistic regression. At this point the results listed in the theoretical section can be used (along with calculating the share of recruits and separations to employment, separation rates, and growth rates for each firm) in conjunction with equation (6) to produce an estimate of the labor supply elasticity facing each firm.⁷

To provide some intuition on the models being estimated, consider the analysis of separations to employment. A large (in absolute value) coefficient on the log earnings variable implies that a small decrease in an individual's earnings will greatly increase the probability of separating in any given period. In a perfectly competitive economy, we would expect this coefficient to be infinitely high. Similarly, a very small coefficient implies that the employer can lower the wage rate without seeing a substantial decline in employment. One concern with this procedure is that this measure of monopsony power is actually proxying

⁷Each equation was also estimated with an indicator variable for whether the employment spell was in progress at the beginning of the data window to correct for potential bias of truncated records. Additionally, all models were reestimated using only job spells for which the entire job spell was observed, with no substantial differences observed between these models.

for high-wage firms, reflecting an efficiency wage view of the economy where firms pay a wage considerably above the market wage in exchange for lower turnover. This is much more of a concern in the full economy estimate of the labor supply elasticity to the firm found elsewhere in the literature than in my firm-level estimation since the models in this paper are run separately by firm. The logic behind this difference is that in the full economy model cross-sectional variation in the level of earnings is used to identify the labor supply elasticity. In a firm-specific model, however, the labor supply elasticity of firm A does not mechanically depend on the level of earnings at firm B. This efficiency wage hypothesis will be directly tested.

Analysis

In addition to the full-economy models of monopsony, I include the concentration ratio and firm-level labor supply elasticity measures in earnings regressions. This provides direct evidence of the effect of firm market power on earnings, a feature not possible in the full-economy models. Additionally, it serves as a test of the efficiency wage hypothesis, which predicts that firms with low estimated labor supply elasticities will pay the highest wages. The main focus of this paper is on this model, explicitly written as:

$$\log(\text{quarterly earnings}_{ij}) = \beta \text{marketpower}_j + \gamma X_{ij} + \delta Y_j + \theta Z_i + \varepsilon_{ij} \quad (18)$$

The dependent variable is the natural log of individual i 's quarterly earnings in employment spell j . The market power variable represents firm j 's estimated labor supply elasticity or the share of the local working population employed at the firm. X is a vector of person and firm characteristics, which may vary by the employment spell, including age, age-squared, tenure (quarters employed at firm), tenure-squared, education⁸, gender, race, ethnicity, year

⁸Reported educational attainment is only available for about 10 percent of the sample, although sophisticated imputations of education are available for the entire sample. The results presented in this paper correspond to the full sample of workers (reported education and imputed education). All models were also run on the sample with no imputed data, and no substantive differences were observed. In particular, since the preferred specification includes person fixed-effects, and thus educational attainment drops out of the

effects, indicator variables for the two-digit NAICS sector, and the size (employment) of the firm. Y is a vector of firm fixed-effects, Z is a vector of person fixed-effects, and ε is the error term. Time-invariant characteristics in X are excluded in models with person or firm fixed-effects.

Finally, to examine whether there is a disproportionate impact of imperfect competition on workers near the bottom of the earnings distribution, I construct a counterfactual earnings distributions in which each firm’s labor supply elasticity is increased. The counterfactual distribution is constructed according to the unconditional quantile approach decomposition suggested in Firpo et al. (2010). Unconditional quantile regression, first introduced in Firpo et al. (2009), estimates the parameters of a regression model as they relate to the quantiles of the dependent variable. This contrasts with traditional quantile regression, which estimates parameters corresponding to the conditional (on the included regressors) quantiles of the dependent variable. The unconditional quantile approach is most advantageous in models with relatively low R-squared (i.e. all wage regressions) since the quantiles of y are most likely to diverge from the quantiles of \hat{y} (predicted dependent variable) in this scenario.

Under this approach, unconditional quantile regressions are performed on every 5th quantile of the earnings distribution using the same model as Equation (18). The estimated coefficients on the labor supply elasticity variable from each regression will then be used to simulate the impact of a one unit increase in the labor supply elasticity to the firm on earnings in the associated quantile.

5 Results

Summary Statistics

Table 1 reports summary statistics from my analysis sample. Since the unit of observation is the employment spell, and only dominant jobs are included, some statistics deviate slightly

model, this is of little concern.

from typical observational studies of the labor market (such as a nearly even split of job spells between men and women). The average employment spell lasts about two and a half years, with more than sixty percent of spells resulting from a move from another job. The quarterly nature of the LEHD data make it difficult to precisely identify⁹ whether an individual separated to employment or nonemployment, and therefore the proportion of separations to employment is slightly higher than comparable statistics reported in Manning (2003).

The average firm in my sample employs nearly 3000 workers and hires almost 500 in a given quarter. Several qualifications must be made for these statistics. First, the distributions are highly skewed, with the median firm employing only 400 and hiring 75 in a given quarter. Second is that statistics are not point in time estimates, but rather totals throughout an entire quarter. Finally, remember that these are at the firm (state-level) rather than at the establishment (individual unit) level. Also of note are the employment concentration ratios, with the average firm employing roughly 9 percent of their county's industry specific labor force.

Location-Based Measure

As previously noted, many studies have attempted to search for evidence of monopsony in the labor market through the use of concentration ratios. While this approach was the best available given prior data constraints, it assumes that monopsony power is derived only from geographical constraints.

Table 2 presents the estimated impact of a ten percentage point increase in the concentration ratio in various specifications of Equation (18). These results suggest that, in general, a firm's geographic dominance does not appear to significantly alter the wage bill it pays. Note

⁹The definition used in this paper requires an individual to have no reported earnings for an entire quarter following an employment spell to be defined as a separation to nonemployment, with all other separations coded as a separation to employment. This definition was chosen because it lead to the most conservative (least monopsonistic) results, although the differences were small. The other methods tried involved imputing the time during the quarter at which employment stopped/started based on a comparison of the earnings reported in the last/first quarter to a quarter in which I know the individual worked the entire quarter.

that when the models are run separately by North American Industry Classification System (NAICS) sector, as depicted in Table 3, there is evidence that firms with high concentration ratios in certain industries (such as the utilities sector) pay slightly lower wage bills. However, the effect sizes are small relative to the observed distribution of concentration ratios. Given the small results, and the fact that the industry-specific effects seem to be centered around zero, it seems plausible to conclude that geographic constraints in the labor market play at most a small role in wage determination for the average worker.

Full-Economy Model

I first compute the average labor supply elasticity to the firm prevailing in the economy by estimating Equations (15)-(17) on a pooled sample of all (dominant) employment spells, and combining the results according to Equation (6). Table 4 presents the output of a several specifications of the full-economy monopsony model. The estimated elasticities range from 0.76 to 0.82 depending on the specification (inclusion of fixed effects, etc.). These elasticities are certainly on the small side, implying that at the average firm a wage cut of one percent would only reduce employment by .8 percent. However, this magnitude is still within the range observed by Manning (2003) in the NLSY79. Additionally, even the inclusion of fixed-effects still puts many more restrictions on the parameter estimates than separate estimations for each firm. Based on a comparison of the full-economy model and the firm-level model presented in the next section, the failure to fully saturate the full economy model likely produces downward biased estimates. A detailed discussion of factors which may attenuate these estimates, as well as structural reasons we should expect these results from US data, is given in the “Discussion and Extensions” section.

Firm-Level Measure

Table 5 presents the elasticities estimated through Equations (15)-(17). The first four columns report the average firm-level elasticities of recruitment from employment and nonem-

ployment, and the separation elasticities to employment and nonemployment respectively. The final column combines these elasticities, along with the calculated shares of separations/recruits to/from employment to obtain the labor supply elasticity. Of note is that the labor supply elasticity does not appear to depend substantially on the regressors included in the model. The first three rows report only the long-run elasticities, while the final row describes the elasticities when each quantity is allowed to vary over time. Not accounting for the time-varying nature of the labor supply elasticity, as has been common in the prior literature, appears to underestimate its magnitude by 20%.

Table 6 displays information about the distribution of firms' labor supply elasticities, and Figure 2 presents a kernel density plot of the market power measure¹⁰. This distribution is constructed by separately estimating Equations (15)-(17) for each firm. While the median supply elasticity (0.75) is close to the estimate from the full-economy model, there appears to be significant variation in the market power possessed by firms. I estimate a mean labor supply elasticity of 1.08, however, there are many firms (about 3 percent of the sample) with labor supply elasticities greater than 5. It appears that while there is a nontrivial fraction of firms whose behavior approximates a highly competitive labor market, the majority of the distribution is characterized by significant frictions. While not surprising, to my knowledge this is the first documentation of the large discrepancy in firms' ability to set the wage.

Table 7 reports average labor supply elasticities broken down by NAICS sector. I find significant variation in these estimates across industries. The manufacturing sector appears to enjoy the least wage-setting power, with a labor supply elasticity of 1.82. As manufacturing is likely the most heavily unionized of all sectors, this result is not surprising. By contrast, firms in the health care (0.78) and administrative support (0.72) sectors seem to wield the greatest wage-setting power. This is consistent with the focus on the healthcare market among economists investigating monopsony power.

The central focus of this paper is presented in Table 8, which estimates various specifi-

¹⁰For confidentiality reasons, the long right tail of the kernel density plot has been suppressed

cations of Equation (18) in order to measure the impact of market power on the earnings distribution. Unconditionally, a one unit increase in the labor supply elasticity increases earnings by .13 log points. Even the specifications with the most detailed controls estimate a strong positive relationship between a firm's labor supply elasticity and the earnings of its workers. These estimates range from an impact of 0.05 log points in the model with person fixed-effects to an impact of 0.18 log points with a full compliment of person and firm effects¹¹. This is an important result for the new monopsony literature, because it rules out the possibility that the dynamic model identification strategy is actually identifying high-wage firms whose employees do not often switch jobs due to the high wages.

There is good reason to believe that the estimates in Table 8 are lower bounds of the true impact of firm market power on earnings. Each labor supply elasticity is a weighted average of many more precisely defined elasticities which would more accurately measure a firm's market power over a particular individual. For example, firms likely face different supply elasticities for every occupation, and potentially different elasticities across race and gender groups. From a measurement error perspective, regressing the log of earnings on the average labor supply elasticity to the firm would attenuate the estimates relative to the ideal scenario where I could separately identify every occupation specific elasticity.

While these results are clear evidence that firms exercise their market power, there is reason to believe that firms are not using the majority of labor market power available to them. Bronfenbrenner (1956) first made this point, arguing that most firms in our economy likely faced upward sloping labor supply curves but that these firms would not pay substantially less than the competitive wage. This could be because firm's choose to maximize some function of profits and other quantities such as public perception and worker happiness.

To test this assertion, we can calculate what the coefficient on labor supply elasticity should be in an economy where firms only maximize profits and the mean labor supply

¹¹All models were also run using the time-invariant long run labor supply elasticity rather than the time varying measure. The results of each model which could be run using this measure (firm effects could not be included) were nearly identical.

elasticity is 1.08. This is done by taking the derivative of the coefficient on the marginal product of labor in Equation (14) and dividing this by the coefficient itself, a formula which simplifies to $\frac{1}{\varepsilon^2 + \varepsilon}$. Evaluating this at a labor supply elasticity of 1.08 implies that if firms were exploiting all of their market power then the coefficient on labor supply elasticity in Table 8 should be about 0.45, roughly 2.5 times greater than the estimated 0.18. Even assuming a high degree of measurement error in the assignment of the average labor supply elasticity to all workers in a firm would likely not account for this disparity. One possibility is that firms reduce labor costs through other avenues than wages which are more easily manipulated such as benefits. Alternatively, this may be evidence that firms do not solely maximize profits, but instead maximize some combination of profits and other quantities (i.e. public perception).

Counterfactual Distribution

Table 9 details the disproportionate effect which firms' market power has on workers at the low end of the earnings distribution. Assuming a one unit increase in the labor supply elasticity for each firm (approximately 1 standard deviation), the 10th percentile of the earnings distribution increases by 0.09 log points under the counterfactual assumption, while the median worker sees an increase of 0.04 log points and the 90th percentile remains unchanged. The nonlinear impacts are also clearly seen in the unconditional quantile regression coefficients, which are 4-5 times greater than the OLS coefficient at lower quantiles and essentially zero at the upper end of the distribution.

Standard measures of inequality are also reported in Table 9 for both the empirical and counterfactual distributions. A one unit increase in firms' labor supply elasticity is associated with a 9 percent reduction in the variance of the earnings distribution (0.94 to 0.86 log points). Similarly, we see decreases in the 90-10 ratio (1.32 to 1.3), 50-10 ratio (1.18 to 1.16), and 90-50 ratio (1.12 to 1.11).

These results could arise from a number of different scenarios, the examination of which is

beyond the scope of the current paper. It may reflect low-ability workers having few outside options for employment. This could be due to strict mobility constraints, a less effective job referral network (Ioannides and Loury, 2004), lower job search “ability” (Black, 1981), or simply being qualified for fewer jobs. Another mechanism through which a firm’s market power might differentially affect low wage workers is gender discrimination, as suggested by Hirsch et al. (2010) or racial discrimination. These questions deserve a much deeper treatment, and should be explored in future research.

Figure 3 plots both the empirical earnings distribution and the counterfactual distribution under a more drastic assumption which more closely approximates perfect competition, that each firm’s labor supply elasticity is increased by a factor of 10 (median elasticity goes from .74 to 7.4). The variance of the counterfactual distribution is considerably lower, with nearly all of the movement occurring in the lower half of the distribution. The striking fact about Figure 3 is that the Burdett and Mortensen model predicts this same behavior as the arrival rate of job offers increases.

It is important to note that the results in the counterfactual distribution is estimated from a model which includes all person and firm controls, but no person or firm fixed effects. This is because identifying off of within person/firm variation in a sense redefines the unconditional quantiles of the distribution, and can introduce substantial bias into the results. Given that the OLS estimates of the impact of firm market power are larger in the specifications which include fixed effects, the results in Table 9 should be taken as lower bounds.

Discussion and Extensions

The labor supply elasticities reported in this paper imply that firms possess a high degree of power in setting the wage. For a variety of reasons, these elasticities are on the lower end of those present in the literature. In this section I address the factors which contribute to these results.

First, it should be noted that the only other studies to estimate the labor supply elasticity

to the firm with comprehensive administrative data used European data. Given the veryis a determinant of earnings inequality restrictive (from the point of view of the employer) employment laws in place in many European countries, this result is not surprising. Assuming that job security accrues over time within firm but drops following a transition to a new firm, any law which makes it more difficult to fire a worker effectively lowers the cost to the employee of switching jobs because job security is less of a factor.

One potential criticism of the labor supply elasticities derived in this paper is that the data do not contain detailed occupation characteristics. This problem is mitigated by the fact that the measures are constructed at the firm level in that I am only comparing workers in the same firm in the construction of a firm's monopsony power. Additionally, previous studies such as Hirsch et al. (2010) and Manning (2003) find that the addition of individual-level variables had little impact on the estimated labor supply elasticities and that it was the addition of firm characteristics which altered the results. As a further check of this problem, I compute the aggregate monopsony measures in the NLSY, as done in Manning (2003), both with and without detailed occupation characteristics. As shown in Table 10, I find that the difference between these labor supply elasticities is about 0.2 and is not statistically significant. Keep in mind that even if this difference were statistically significant, the estimates in this paper are still a long way from implying perfect competition. Thus, I conclude that the absence of occupation controls in the LEHD data will not seriously bias the results of this study.

A potentially more serious problem in the estimation of the labor supply elasticity to the firm is endogenous mobility. Consider the standard search theory model with on the job search: A worker will leave their current job if they receive a higher wage offer from another firm. Their wage at the new firm is then endogenously determined since in effect it was drawn from a distribution truncated at the wage of the their previous job. In this sense, the earnings data for those individuals who were hired away from another job is biased upward, which will bias estimates of the labor supply elasticity to the firm downward. I deal with

the endogenous mobility bias in several different ways. First, I estimate the average earnings premium an individual gets from moving to their n th job (where n is the job number in a string of consecutive employment spells). For instance, workers' earnings increase on average .19 log points when they move from their first to their second jobs. I then reduce the earnings of all job movers by the average premium associated with a move from job $n-1$ to n . For example, all workers in their second jobs of a string of employment spells would have their earnings reduced by .19 log points.¹² The rationale behind this adjustment is that I only observe workers moving from one job to another if they receive a higher wage offer (This is a typical assumption of on-the-job search models, and is overwhelmingly true in the data). Thus, the earnings I observe in the second job are endogenously determined, since they were in a sense drawn from a strictly positive offer distribution.

Second, I recalculate the labor supply elasticities with a Heckman selection correction. In this model I define the selected group as those who separate from one job to another, and use the number of new jobs in an individual's state and industry as the excluded variable. The logic behind this restriction is that the state-industry specific labor market should be highly correlated with the likelihood that an individual moves to a new job, but should be uncorrelated with that individual's unobserved "ability" to move. The inverse Mills ratio from the Heckman selection model is included as a regressor in each of the Equations (15)-(17). As noted in Table 10, each of these corrections leads to a trivial change in the labor supply elasticity distribution.

One final concern regarding endogenous mobility is that we do not observe the complete history of workers, only that within the time-frame of the LEHD infrastructure. Thus, any employment spells in progress at the beginning of our window which are the result of a hire from another firm may introduce bias into the results. To assess the degree to which this is a problem, I again employ the NLSY79. I use a Monte Carlo approach to compare

¹²Define a string of employment spells as consecutive jobs an individual holds with no time spent outside the labor force. In other words, each job transition in a string of employment spells is defined as being a separation to, or recruitment from, employment. An observation takes a default value of 1, 2 if the employment spell is the second in a string of spells, etc.

the estimated labor supply elasticities using the complete worker histories and using only employment spells which occurred in the final third of the sample window. This is the ideal comparison, where the first calculation takes into account the entire work histories of each individual and the second calculation uses only those spells observed after an arbitrary date. The Monte Carlo analysis finds that using the complete worker histories leads to a statistically insignificant decrease of the estimated labor supply elasticity. This implies that the use of some partial histories in this study is not likely a problem, and at worst yields an underestimate of monopsony power.

For the reasons mentioned in this section and probably many others, critics may claim that this paper does not accurately estimate the labor supply elasticity to the firm, and they could be right. As with any identification strategy, this study relies on assumptions, not all of which are testable. But while the average firm's labor supply elasticity may not be exactly 1.08, the variable which I call a supply elasticity is certainly some kind of weighted average highly correlated with mobility and individuals' responsiveness to changes in earnings. The fact that this measure is highly correlated with earnings, especially for those at the bottom of the distribution, tells us that our economy is less competitive than we commonly assume.

6 Conclusion

This study finds evidence of significant frictions in the US labor market, although the severity of these frictions varies greatly between labor markets. I estimate the average firm's labor supply elasticity to be quite monopsonistic at 1.08, however there is a nontrivial fraction of firms who do appear to be operating in an approximately competitive labor market. While identifying the precise frictions which contribute to firms' labor market power is beyond the scope of this study, I can conclude that a firm's geographical dominance alone does not account for all or even most of their ability to affect the wage offer distribution.

I extend the dynamic model-based empirical strategy proposed by Manning (2003) to

identify firm level labor supply elasticities. The use of these measures of firm market power in earnings regressions provides the first direct test of the validity of the new monopsony model. I find that a one unit increase in a firm's labor supply elasticity is associated with a 5-18 percent increase in earnings on average. Further exploring the earnings distribution, I find highly nonlinear effects implying that the negative effects of monopsony power are concentrated at the lower end of the distribution. While these effects are certainly not trivial, it is important to note that there is evidence that firms only utilize a fraction of their market power.

The development of the firm-level measures of labor market power described in this paper could have a significant impact on how we view the interaction of imperfect competition with traditional models of the labor market. Future research will examine topics such as gender/race wage gaps, minimum wage laws, unionization, labor demand over the business cycle, agglomeration, and many others.

References

- J. Abowd and F. Kramarz, "The costs of hiring and separations," *Labour Economics*, vol. 10, pp. 499–530, 2003.
- J. Abowd and L. Vilhuber, "National estimates of gross employment and job flows from the quarterly workforce indicators with demographic and industry detail," *Journal of Econometrics*, vol. 161, pp. 82–99, 2011.
- J. Abowd, F. Kramarz, and D. Margolis, "High wage workers and high wage firms," *Econometrica*, vol. 67, pp. 251–335, 1999.
- J. Abowd, B. Stephens, L. Vilhuber, F. Andersson, K. McKinney, M. Roemer, and S. Woodcock, "The lehd infrastructure files and the creation of the quarterly workforce indicators,"

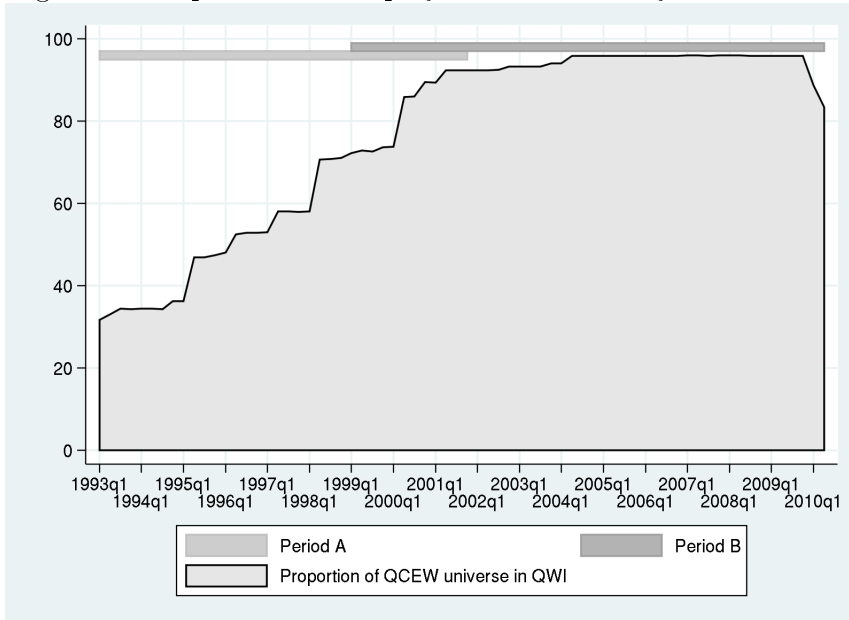
- in *Producer Dynamics: New Evidence from Micro Data*, J. B. J. T. Dunne and M. J. Roberts, Eds. The University of Chicago Press, 2009, pp. 149–234.
- K. Adamache and F. Sloan, “Unions and hospitals, some unresolved issues,” *Journal of Health Economics*, vol. 1(1), pp. 81–108, 1982.
- M. Black, “An empirical test of the theory of on-the-job search,” *Journal of Human Resources*, vol. 16(1), pp. 129–140, 1981.
- M. Bronfenbrenner, “Potential monopsony in labor markets,” *Industrial and Labor Relations Review*, vol. 9(4), pp. 577–588, 1956.
- P. Brummund, “Variation in monopsonistic behavior across establishments: Evidence from the Indonesian labor market,” 2011, working Paper.
- R. Bunting, *Employer Concentration in Local Labor Markets*. University of North Carolina Press, 1962.
- K. Burdett and D. Mortensen, “Wage differentials, employer size, and unemployment,” *International Economic Review*, vol. 39(2), pp. 257–273, 1998.
- D. Card and A. Krueger, *Myth and Measurement: The New Economics of the Minimum Wage*. Princeton University Press, 1995.
- B. Depew and T. Sorensen, “Elasticity of labor supply to the firm over the business cycle,” 2011, working Paper.
- R. Feldman and R. Scheffler, “The union impact on hospital wages and fringe benefits,” *Industrial and Labor Relations Review*, vol. 35, pp. 196–206, 1982.
- S. Firpo, N. Fortin, and T. Lemieux, “Unconditional quantile regressions,” *Econometrica*, vol. 77(3), pp. 953–973, 2009.

- , *Decomposition Methods in Economics*, 4th ed., ser. Handbook of Labor Economics. Elsevier, 2010, pp. 1–102.
- Y. Genda and A. Kondo, “Long-term effects of a recession at labor market entry in japan and the united states,” *Journal of Human Resources*, vol. 45(1), pp. 157–196, 2010.
- E. Goux and E. Maurin, “Persistence of inter industry wage differentials: A reexamination using matched worker-firm panel data,” *Journal of Labor Economics*, vol. 17, pp. 492–533, 1999.
- B. Hirsch and E. Schumacher, “Monopsony power and relative wages in the labor market for nurses,” *Journal of Health Economics*, vol. 14(4), pp. 443–476, 1995.
- B. Hirsch, T. Schank, and C. Schnabel, “Differences in labor supply to monopsonistic firms and the gender pay gap: An empirical analysis using linked employer-employee data from germany,” *Journal of Labor Economics*, vol. 28(2), pp. 291–330, 2010.
- R. Hofler and K. Murphy, “Underpaid and overworked: Measuring the effect of imperfect information on wages,” *Economic Inquiry*, vol. 30(3), pp. 511–529, 1992.
- R. Hurd, “Equilibrium vacancies in a labor market dominated by non-profit firms: The shortage of nurses,” *Review of Economics and Statistics*, vol. 55(2), pp. 234–240, 1973.
- Y. Ioannides and L. Loury, “Job information networks, neighborhood effects, and inequality,” *Journal of Economic Literature*, vol. 42(4), pp. 1056–1093, 2004.
- L. Jacobson, R. LaLonde, and D. Sullivan, “Earnings losses of displaced workers,” *American Economic Review of Economics and Statistics*, vol. 83(3), pp. 685–709, 1993.
- L. Kahn, “The long-term consequences of graduating from college in a bad economy,” *Labour Economics*, vol. 17(2), pp. 303–316, 2010.
- P. Kuhn, “Is monopsony the right way to model labor makrets,” *International Journal of the Economics of Business*, vol. 11(3), pp. 369–378, 2004.

- C. Link and J. Landon, "Monopsony and union power in the market for nurses," *Southern Economic Journal*, vol. 41, pp. 649–659, 1975.
- C. Link and R. Settle, "Labor supply responses of married professional nurses: New evidence," *Journal of Human Resources*, vol. 14, pp. 256–266, 1979.
- A. Manning, *Monopsony In Motion*. Princeton University Press, 2003.
- , *Imperfect Competition in the Labor Market*, 4th ed., ser. Handbook of Labor Economics. Elsevier, 2010, pp. 973–1041.
- J. McCall, "Economics of information and job search," *Quarterly Journal of Economics*, vol. 81(1), pp. 113–126, 1970.
- D. Mortensen, *Wage Dispersion*. MIT Press, 2003.
- W. Oi, "Labor as a quasi-fixed factor," *Journal of Political Economy*, vol. 70, pp. 538–555, 1962.
- P. Oyer, "Initial labor market conditions and long-term outcomes for economists," *The Journal of Economic Perspectives*, vol. 20(3), pp. 143–160, 2006.
- , "The making of an investment banker: Macroeconomic shocks, career choice and lifetime income," *The Journal of Finance*, vol. 63(12), pp. 2601–2628, 2008.
- A. Pigou, *The Economics of Welfare*. Macmillan, 1924.
- S. Polachek and J. Robst, "Employee labor market information: comparing direct world of work measures of workers' knowledge to stochastic frontier estimates," *Labour Economics*, vol. 5(2), pp. 231–242, 1998.
- M. Ransom and R. Oaxaca, "New market power models and sex differences in pay," *Journal of Labor Economics*, vol. 28(2), pp. 267–290, 2010.

- M. Ransom and D. Sims, "Estimating the firm's labor supply curve in a 'new monopsony' framework: School teachers in missouri," *Journal of Labor Economics*, vol. 28(2), pp. 331–355, 2010.
- J. Robinson, *The economics of imperfect competition*. Macmillan, 1933.
- R. Rogerson, R. Shimer, and R. Wright, "Search-theoretic models of the labor market: A survey," *Journal of Economic Literature*, vol. 43(4), pp. 959–988, 2005.
- J. Schmieder, "Labor costs and the evolution of new establishments," 2010, working Paper.
- R. Shimer, "The assignment of workers to jobs in an economy with coordination frictions," *Journal of Political Economy*, vol. 113(5), pp. 996–1025, 2005.
- G. Stigler, "Information in the labor market," *Journal of Political Economy*, vol. 70(5), pp. 94–104, 1962.
- D. Sullivan, "Monopsony power in the market for nurses," *Journal of Law and Economics*, vol. 32, pp. S135–S178, 1989.
- E. Wasmer, "General versus specific skills in labor markets with search frictions and firing costs," *American Economic Review*, vol. 96(3), pp. 811–831, 2006.

Figure 1: Proportion of Employment Covered by the LEHD Infrastructure



Reproduced with permission from Abowd and Vilhuber (2011)

Figure 2: Distribution of Labor Supply Elasticities

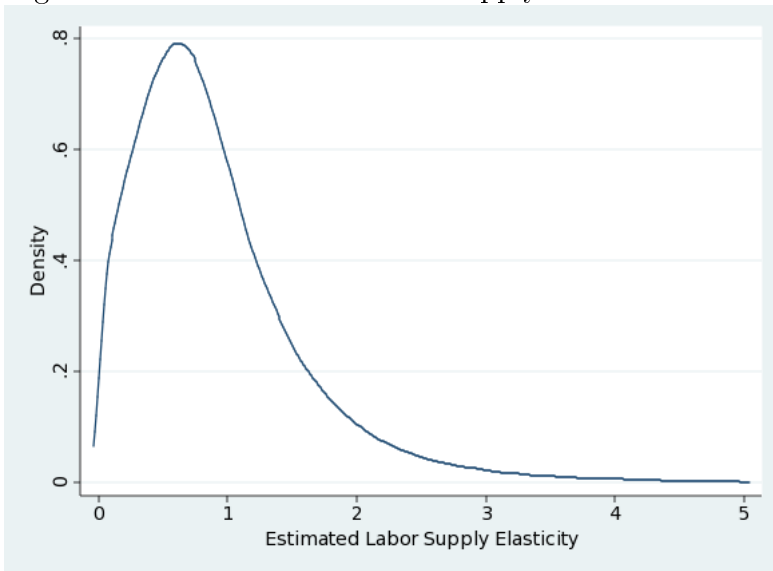


Figure 3: Empirical and Counterfactual Distributions




Table 1: Summary Statistics

Variable	Mean	Std Dev
Age	38	15.2
Female	0.5	0.5
White	0.77	0.42
Hispanic	0.14	0.34
< High School	0.14	0.34
High School Diploma	0.29	0.45
Some College	0.32	0.47
College Degree+	0.25	0.43
Tenure (Quarters)	10.1	10.7
Log(Quarterly Earnings)	8.5	1
Firm Concentration	0.01	0.02
Firm Industry-Concentration	0.09	0.16
Firm Hires per Quarter	493	1592
Firm Employment	2962	10772
Separation Rate	0.15	0.15
Employment Growth Rate	1.01	0.15
Recruited from Employment	0.64	0.48

Table 2: Impact of Firm Concentration on Earnings

Impact of a ten percentage point increase in concentration ratio on log(earnings)	0.0213	0.0053	0.0109	0.0066	0.0114
Demographic and human capital controls	No	Yes	Yes	Yes	Yes
Employer controls	No	No	Yes	Yes	Yes
Tenure Controls	No	No	No	Yes	Yes
State fixed-effects	No	No	No	No	Yes
R-Squared	0.0013	0.2369	0.3300	0.3438	0.3502
Observations	325,630,000	325,630,000	325,630,000	325,630,000	325,630,000

*A pooled national sample of all dominant employment spells is used in this set of regressions. The dependent variable is the natural log of quarterly earnings. Demographic and human capital controls include: age, age-squared, and indicator variables for gender, ethnicity, racial status, and education level. Employer controls include indicator variables for each of the 20 NAICS sectors and number of employees working at the firm. Tenure controls include the length (in quarters) of the employment spell, as well as its squared term. Year effects are included in all models. The results are not reported for the models with firm and person fixed effects because the coefficient was deemed to be biased due to severe multicollinearity. Standard errors are not reported because all t-statistics are greater than 50. Observation counts are rounded to the nearest 10,000 for confidentiality reasons.

Table 3: Concentration Ratio Regressions by Industry	
Industry	Impact of a ten percentage point increase in concentration ratio on log earnings
Agriculture	0.0055
Mining/Oil/Natural Gas	0.0071
Utilities	-0.0760
Construction	-0.0157
Manufacturing	0.0050
Wholesale Trade	-0.0142
Resale Trade	-0.0009
Transportation	0.0361
Information	-0.0308
Finance and Insurance	-0.015
Real Estate and Rental	0.022
Profession/Scientific/Technical Services	0.019
Management of Companies	0.056
Administrative Support	-0.01
Educational Services	-0.005
Health Care and Social Assistance	0.016
Arts and Entertainment	0.046
Accommodation and Food Services	0.021
Other Services	-0.129
Public Administration	-0.013

*A pooled national sample of all dominant employment spells is used in this set of regressions. The dependent variable is the natural log of quarterly earnings. Demographic and human capital controls include: age, age-squared, and indicator variables for gender, ethnicity, racial status, and education level. Employer controls include the number of employees working at the firm. Tenure controls include the length (in quarters) of the employment spell, as well as its squared term. Year effects are included in all models.

Table 4: Full-Economy Estimate of the Labor Supply Elasticity to the Firm		
Full sample	Full sample with firm FE	Only firms with an individually estimated elasticity
.76	.82	.81

*These labor supply elasticities were obtained by estimating equations (15)-(17), on a pooled sample of all (dominant) employment spells. Each model contained age, age-squared, along with indicator variables for female, nonwhite, Hispanic, high school diploma, some college, college degree or greater, year, and each of 20 NAICS sectors.

Model	ε_R^E	ε_R^N	ε_S^E	ε_S^N	ε
Earnings Only	0.41	0.1	-0.41	-0.5	0.84
No Education Controls	0.43	0.3	-0.43	-0.52	0.89
Full Model	0.47	0.46	-0.47	-0.54	0.95
Full Model (Time-Varying)	0.6	0.59	-0.6	-0.67	1.08

The first row represents estimates from equations (15)-(17) where the only regressor in each model is log earnings. The second row estimates the same equations, and includes age, age-squared, along with indicator variables for female, nonwhite, Hispanic, and year effects. The third row adds indicator variables for completing a high school diploma, some college, and college degree or greater. The first four columns report the average firm-level elasticities of recruitment from employment and nonemployment, and the separation elasticities to employment and nonemployment respectively. The final column combines these elasticities, along with the calculated shares of separations/recruits to/from employment, separation rates, and growth rates to obtain the labor supply elasticity. The first three rows report only the long-run elasticities, while the fourth row describes the elasticities when a steady-state is not assumed, and they are allowed to vary over time.

Table 6: Distribution of Estimated Firm-Level Labor Supply Elasticities					
	Percentiles				
Mean	10th	25th	50th	75th	90th
1.08	0.22	0.44	0.75	1.13	1.73

*Three separate regressions, corresponding to equations (15)-(17), were estimated separately for each firm in the data which met the conditions described in the data section. The coefficients on log earnings in each regression were combined, weighted by the share of recruits and separations to employment, separation rates, and growth rates according to equation (6) to obtain the estimate of the labor supply elasticity to the firm. Demographic and human capital controls include: age, age-squared, and indicator variables for gender, ethnicity, racial status, and education level. Employer controls include number of employees working at the firm and industry indicator variables. Year effects are included in all models.

Table 7: Mean Labor Supply Elasticity by Sector	
NAICS Sector	Mean Labor Supply Elasticity
Agriculture	1.43
Mining/Oil/Natural Gas	1.52
Utilities	1.18
Construction	1.42
Manufacturing	1.82
Wholesale Trade	1.48
Resale Trade	1.03
Transportation	1.47
Information	1.17
Finance and Insurance	1.27
Real Estate and Rental	1.01
Profession/Scientific/Technical	1.17
Services	
Management of Companies	1.17
Administrative Support	0.72
Educational Services	0.91
Health Care and Social Assistance	0.78
Arts and Entertainment	0.94
Accommodation and Food Services	0.85
Other Services	1.04
Public Administration	1.19

*The numbers in this table represent averages by NAICS sector of the estimated labor supply elasticity to the firm. Three separate regressions, corresponding to equations (15)-(17), were estimated separately for each firm in the data which met the conditions described in the data section. The coefficients on log earnings in each regression were combined, weighted by the share of recruits and separations to employment, separation rates, and growth rates according to equation (6) to obtain the estimate of the labor supply elasticity to the firm. Demographic and human capital controls include: age, age-squared, and indicator variables for gender, ethnicity, racial status, and education level. Employer controls include number of employees working at the firm. Year effects are included in all models.

	0.13	0.11	0.05	0.03	0.03	0.05	0.09	0.18
Coefficient on labor supply elasticity								
Demographic controls	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Employer controls	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Tenure controls	No	No	No	Yes	Yes	Yes	Yes	Yes
State fixed-effects	No	No	No	No	Yes	Yes	Yes	Yes
Person fixed-effects	No	No	No	No	No	Yes	No	Yes
Firm fixed-effects	No	No	No	No	No	No	Yes	Yes
R-Squared	0.005	0.238	0.312	0.331	0.338	0.784	0.90	0.99

*A pooled national sample of all dominant employment spells subject to the sample restriction described in the data section is used in this set of regressions. The dependent variable is the natural log of quarterly earnings. Demographic controls include: age, age-squared, and indicator variables for gender, ethnicity, racial status, and education level. Employer controls include the number of employees working at the firm and industry indicator variables. Tenure controls include the length (in quarters) of the employment spell, as well as its squared term. Year effects are included in all models. These results are unweighted, however all models were also estimated with demographic weights constructed by the author. There were no significant differences between the weighted and unweighted models. Standard errors are not reported because the t-statistics range from 500-1000, but are available upon request along with all other estimated coefficients. There are 267,310,000 observations in each specification.

Table 9: Counterfactual Distribution Analysis

Change (log points) in Quantiles of the Earnings Distribution					
Quantile	10th	25th	50th	75th	90th
Change in log(earnings)	0.09	0.05	0.04	0.01	0.00
Inequality measure	Variance	90-10	50-10	90-50	
Earnings distribution	.94	1.32	1.18	1.12	
Counterfactual distribution	.86	1.30	1.16	1.11	

*The counterfactual distribution was constructed by estimating unconditional quantile regressions at every fifth quantile of the earnings distribution, and using the supply elasticity coefficient from each regression to simulate the effect at each quantile of a one-unit increase of the labor supply elasticity. Demographic and human capital controls include: age, age-squared, and indicator variables for gender, ethnicity, racial status, and education level. Employer controls include the number of employees working at the firm and industry indicator variables. Tenure controls include the length (in quarters) of the employment spell, as well as its squared term. Year effects are included in all models.

Table 10: Robustness Checks

*Panel A: NLSY comparisons	With versus without occupational effects	Full history versus partial history	
Bootstrapped difference in labor supply elasticity	0.20	-.46	
Std Error	0.14	.76	
**Panel B: Endogenous mobility corrections	Uncorrected labor supply elasticity	Earnings of job changers adjusted downward	Control for Heckman selection correction
Median of distribution	.75	.74	.76

*Panel A: Equations (15)-(17) were estimated on a sample of employment spells from the NLSY79 from 1979-1996 (the last year for which detailed information on recruitment and separation dates are available). The specifications include the same variables available through the LEHD data: age, age-squared, year effects, along with gender, ethnicity, race, industry, and education indicators. The first column compares the labor supply elasticities with and without the inclusion of occupational fixed effects. The second column compares the labor supply elasticities with and without the assumption that only the last third of every individual's work history is known.

**Panel B: The second column represents a recalculation of the labor supply elasticity in which workers who are recruited away from another job have their earnings adjusted downward by the average premium of moving from job n to job $n+1$. The third column represents a recalculation of the labor supply elasticity in which the inverse Mills ratio of a Heckman selection model for mobility is controlled for in each of Equations (15)-(17). The omitted category in the Heckman model is the number of new local jobs in each workers current industry.