

Webber, Douglas A.

Working Paper

Firm-level monopsony and the gender pay gap

IZA Discussion Papers, No. 7343

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: Webber, Douglas A. (2013) : Firm-level monopsony and the gender pay gap, IZA Discussion Papers, No. 7343, Institute for the Study of Labor (IZA), Bonn

This Version is available at:

<https://hdl.handle.net/10419/71636>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

IZA DP No. 7343

Firm-Level Monopsony and the Gender Pay Gap

Douglas A. Webber

April 2013

Firm-Level Monopsony and the Gender Pay Gap

Douglas A. Webber

*Temple University
and IZA*

Discussion Paper No. 7343
April 2013

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0
Fax: +49-228-3894-180
E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Firm-Level Monopsony and the Gender Pay Gap^{*}

Using a dynamic labor supply model and linked employer-employee data, I find evidence of substantial search frictions, with females facing a higher level of frictions than males. However, the majority of the gender gap in labor supply elasticities is driven by across firm sorting rather than within firm differences, a feature predicted in the search theory literature, but which has not been previously documented. The gender differential in supply elasticities leads to 3.3% lower earnings for women. Roughly 60% of the elasticity differential can be explained by marriage and children penalties faced by women but not men.

JEL Classification: J42, J71

Keywords: monopsony, discrimination

Corresponding author:

Doug Webber
Temple University
Ritter Annex 883
1301 Cecil B. Moore Ave.
Philadelphia, PA 19122
USA
E-mail: douglas.webber@temple.edu

^{*} 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, and Kevin Hallock. I would also like to sincerely thank J. Catherine Maclean, Erika McEntarfer, Ben Ost, and Michael Strain for their many helpful comments. All errors are mine.

1 Introduction

The male-female wage gap has long been a fixture of the labor economics literature (see Altonji and Blank (1999), Blau and Kahn (2008), or Bertrand (2011) for excellent summaries). While certainly not true of all studies, an abundance of the literature evaluates factors which contribute to the gap by (1) estimating wage equations controlling for all observable characteristics between men and women, (2) adjusting for differences in the observables through a decomposition method, and (3) interpreting all or some of the remaining gap as discrimination or some other unobservable factor. The interaction between the model coefficients and the group level differences in each observable variable is taken to be the contribution of that variable to the wage gap. This is a perfectly reasonable strategy, and in effect is exactly what this study does.

The difference between this study and the previous literature is the ability to control for detailed firm-level measures of labor market power. An assumption of most of the literature, dating back to Becker (1971), is that the structural features of the labor market are the same for both men and women. By this I mean that if we could perfectly control for all ability-related personal characteristics then two workers at the same firm doing the same job must be paid the same wage, and if not then the residual difference is due to discrimination. Becker's analysis is underlied by the belief that the perfectly competitive market forces would drive discriminating employers out of the labor market in the long run.

Recent evidence refutes these assumptions, finding significant frictions in the labor market Manning (2003); Webber (2012), as well as theoretical Bowlus (1997) and empirical (Ransom and Oaxaca, 2010; Hirsch et al., 2010) evidence that these frictions may differ by gender. This implies that firm characteristics may play a large role in wage determination, and that firm fixed-effects would not be enough to explain the effect of these frictions on the wage gap. Additionally, it implies that part of the wage gap might be explainable through firms' profit maximization; in other words, price discrimination rather than taste discrimination.

Using the Longitudinal Employer Household Dynamics (LEHD) infrastructure, linked

employer-employee data from the US Census Bureau, I separately estimate the labor supply elasticities for men and women at nearly 100,000 firms spanning 47 states. This strategy allows me to evaluate two distinct avenues through which firm market power may contribute to the male-female wage gap, within firm and across firm disparities in the gender specific labor supply elasticities. To understand the difference, consider an economy where the male labor supply elasticity is greater (more competitive) than the female labor supply elasticity by the same magnitude at every firm. Now consider a parallel economy with the same aggregate difference in gender specific labor supply elasticities, but instead of each firm having the same differential there is no difference in the elasticities at any firm but instead women disproportionately work at firms with low labor supply elasticities.

In both economies the market-level labor supply elasticities for men and women and the implied impact of market power on the wage gap are the same, but the mechanisms are quite different. In the first economy women face less competition for their labor (potential mechanisms will be discussed later), a fact which is exploited by firms in the form of lower wages for equally qualified workers. In the second economy, each firm pays its workers the same wage rate regardless of gender, with the difference in market-level wages arising from segmentation of the labor force, with male-dominated firms operating in more competitive labor markets than multi-sex firms. Note in this second economy traditional discrimination is still very possible, but it operates through the employment margin rather than the wage margin.

Using a dynamic labor supply model to separately estimate male and female labor supply elasticities for each firm in my sample, I find strong evidence of across-firm labor supply elasticity differentials, but only small within-firm differentials. At firms where I am able to estimate both a male and female elasticity, I find average (worker-weighted) labor supply elasticities of 0.98 and 0.94 for men and women respectively. However, the average labor supply elasticities are 1.09 and 0.94 for men and women respectively when I examine firms for which I can estimate at least one of the gender-specific elasticities. Furthermore, I can

directly estimate the impact of the gender gap in search frictions on the male-female earnings gap. I estimate that on average gender-specific search frictions lead to 3.3% lower earnings for women relative to men.

This paper contributes to the current literature in several important ways. First, the previous literature has only been able to examine how labor supply elasticities differ by gender at the market level. Thus, this literature can only produce two market-level elasticities (one male and one female), whereas the current paper produces firm-specific elasticities for more than one hundred thousand firms. This allows me to characterize the distribution and composition of the gender labor supply elasticity gap (within versus across firm, industry, etc.). Second, when evaluating the impact of imperfect competition on the gender wage gap, the previous literature has been forced to provide a theoretically implied impact (because two market-level elasticities cannot be used in statistical inference) rather than a directly estimated impact as is done in this study. I find that the theoretically implied impact drastically overstates the directly estimated impact. Finally, the model used in this study is considerably more flexible than the gender pay gap models which have previously been estimated, allowing for substantially more firm heterogeneity as well as allowing the labor supply elasticity to vary over time. Furthermore, I find evidence that women face mobility penalties for marriage and children which men do not face, which explains more than half of the gender labor supply elasticity differential.

The paper is organized as follows, Section 2 describes motivation behind looking at the gender wage gap through a monopsony perspective. Section 3 discusses the previous literature. Section 4 lays out the theoretical foundation for this study. The data and methods are described in Section 5. Section 6 presents the results and sensitivity analyses, and Section 7 concludes.

2 Monopsony and the Gender Pay Gap

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 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, oligopsony, 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.

Webber (2012) discusses in detail some of the many potential sources of a firm’s monopsony power, including: geographic constraints, moving costs, firm specific human capital, job security, asymmetric information, compensating differentials, and more. In this study, I will focus instead on factors which may cause a difference in the labor supply elasticities for men and women.

Many of the factors which may cause a difference in the male and female labor supply elasticities are sociological in nature. For example, on average the male’s job within a marriage is the dominant job. So a family may make locational decisions based primarily on the job prospects for the husband, thus forcing the wife to search for a job only in a local labor market centered around her husband’s place of employment. Women may also place a greater importance on non-wage benefits offered by employers, such as flexible work schedules or other family-friendly practices which limit the number of jobs which are suitable. For instance, if female workers are more risk averse, in terms of either job or earnings stability, than their male counterparts, then this may act as a compensating differential which would

manifest itself in the form of a lower labor supply elasticity (Bonin et al., 2007). Additionally, since the core cause of a firm's monopsony power lies in the fact that workers do not have an infinite stream of job offers, discrimination in the hiring process against women would lead to a lower labor supply elasticity (and thus lower wages) even for women at nondiscriminating firms because they would have fewer outside options. This is an important point which is explored within the context of an equilibrium search model in Black (1995).

Much of the recent labor literature views monopsony power through a search theory context, a framework which has also been used to model gender wage differentials. Bowlus (1997) extends the standard on the job search model to allow for an individual to occupy one of three states (employment, unemployment, and nonparticipation), and allows the underlying search parameters to vary by gender. A structural estimation of this model using the National Longitudinal Survey of Youth (NLSY) in 1979 concludes that the search behavior of men and women is statistically different, with women facing a lower arrival rate of job offers and having a higher separation rate than men. Bowlus (1997) finds that this difference in search behavior explains between 20 and 30 percent of the gender wage gap. Furthermore, Bowlus (1997) concludes that this differential would likely manifest itself through firm segmentation by gender rather, with women more likely to work in low wage firms due to the difference in search behavior rather than within-firm differences in pay.

3 Previous Literature

The empirical monopsony literature dates back to Bunting (1962), 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 this literature has been that of nurses in rural 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 elasticity¹. Studying the market for nurses, Sullivan (1989) finds evidence of monopsony using a semistructural 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 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.

Manning (2003) formalized a method for identifying the labor supply elasticity facing the firm off of job to job transitions. This dynamic model of labor supply, which derives its roots from Card and Krueger (1995) and the Burdett and Mortensen (1998) equilibrium search model, is the basis for the model used in this paper. Applying the model 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) uses a novel structural production function approach, and finds strong evidence of monopsony in Indonesian labor markets, estimating labor supply elasticities between 0.6 and 1.0.

A dynamic model of labor supply approach has also been used to evaluate the link between monopsony and the gender pay gap. Two careful studies, Ransom and Oaxaca (2010) and Hirsch et al. (2010) both separately estimate the labor supply elasticities to the firm at the market level 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

¹The 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.

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. These studies conclude that at least one third of the wage gap between men and women can be attributed to firm-level monopsony. It is important to note that this cannot be directly tested in the data used in these studies, but rather is theoretically implied by the difference in gender-specific elasticities at the market level. It should be noted that the proposed link between the gender pay gap and monopsony is not a new idea in the labor literature, with Madden (1973) devoting an entire book to this topic.

The closest analogue to this study in terms of method and data is Webber (2012), which uses linked employer employee data from the U.S. Census Bureau and an extended dynamic labor supply model to study firm-level monopsony.

4 Theoretical Model

Equilibrium search models are the theoretical basis underlying most of the recent monopsony literature. The seminal model of an economy with search frictions is that of Burdett and Mortensen (1998). They develop a model of the economy with on-the-job search 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). Also part of the Burdett-Mortensen class of search models, and of particular relevance to the present study, is Bowlus (1997). The Bowlus model allows for individuals to be out of the labor force and not be search for a job (nonparticipation) in addition to the standard employed and unemployed states. While I do not explicitly estimate either the Burdett and Mortensen or the Bowlus models in this paper, the intuition of monopsony power derived from search frictions is central to this study. The following is a description of the dynamic labor supply model which I estimate.

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:

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

$$s(w) = \delta + \lambda(1 - F(w)) \quad (2)$$

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:

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

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:

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

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. This study estimates the above parameters separately by gender, thus yielding gender-specific labor supply elasticities to the firm for every available firm.

The model presented above implies that, even in a world where all firms are identical and individuals possess equal ability, a difference in the job offer arrival rate across gender will lead to a gender wage gap. This is true even for firms who do not discriminate in a taste-based sense. 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:

$$\underset{w}{Max} \Pi = pQ(L_M) - w_M L_M(w_M) + pQ(L_F) - w_F L_F(w_F) \quad (13)$$

P is the price of the output produced according to the production function Q. The choice of wage w determines the male and female labor supplied to the firm L_M and L_F respectively. Taking first order conditions, substituting $\varepsilon = \frac{w}{L(w)} \frac{\partial L(w)}{\partial w}$, and solving for the gender-specific wage yields:

$$w_M = \frac{pQ'(L_M)}{1 + \frac{1}{\varepsilon_M}} \quad (14)$$

$$w_F = \frac{pQ'(L_F)}{1 + \frac{1}{\varepsilon_F}} \quad (15)$$

The numerator in Equation (14) is simply the marginal product of labor, and ε_M and ε_F

are the gender-specific labor supply elasticities 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 $0 < \varepsilon < \infty$.

5 Data and Methodology

Data

This study uses linked employer-employee data from the U.S. Census Bureau to estimate the gender-specific firm level labor supply elasticities. 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 1990 and 2008². I make several sample restrictions in an attempt to obtain the most economically meaningful results. 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

²The 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.

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 quarter³. I also remove all spells which span fewer than three quarters.⁴ 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 bottom 1 percent of observations

Additionally, I limit the analysis to firms with at least 100 total employment spells of any length over the lifespan of the firm, and 25 employment spells in each estimating equation. After making these restrictions, I am left with two samples of interest, All workers for whom I can estimate a gender-specific labor supply elasticity, and workers who work at firms where I can identify both a male and female elasticity. The first sample is made up of roughly 242 million employment spells, belonging to about 105 million unique individuals, who work at approximately 250 thousand separate firms. The sample requiring each firm to have both a male and female labor supply elasticity has roughly 183 million employment spells, belonging to about 84 million unique individuals, who work at approximately 100 thousand separate firms.

³This 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.

⁴The relaxation of this assumption does not appreciably alter any of the reported results.

Empirical Strategy

The construction of the labor supply elasticities presented in this paper most closely represents an augmented gender-by-firm level implementation of the methodology proposed in Manning (2003), with the extension allowing for a time-varying elasticity described above.

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 recruitment share, separation share, growth rate, and separation rate for each firm. Each of the following models is run separately by gender for every firm in the sample, where the unit of observation is an employment spell. Looking first at the separation elasticities, I model separations to nonemployment as a Cox proportional hazard model given by

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

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 race, age, education, firm size, and year control variables (time-invariant firm characteristics such as industry cannot be included because the model is run at the firm level). 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

$$\lambda^E(t|\beta^{E,sep}\log(earnings)_i + X_i\gamma^{E,sep}) = \lambda_0(t) \exp(\beta^{E,sep}\log(earnings)_i + X_i\gamma^{E,sep}) \quad (17)$$

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

$$P_{rec} = \frac{\exp(\beta^{E,rec} \log(\text{earnings})_i + X_i \gamma^{E,rec})}{1 + \exp(\beta^{E,rec} \log(\text{earnings})_i + X_i \gamma^{E,rec})} \quad (18)$$

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.⁵

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 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 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 directly testable, and is rejected as an explanation later in the paper.

⁵Each 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.

Analysis

In order to directly estimate the impact of firm-level monopsony on the gender pay gap, we must estimate two quantities: the male-female gap in labor supply elasticities and the impact of the labor supply elasticity on earnings. The elasticity gap can be derived from the above results. The impact of the labor supply elasticity on earnings can be estimated from the following equation.

$$\log(\text{quarterly earnings}_{ij}) = \beta \text{elasticity}_{jg} + \gamma X_{ij} + \delta Y_j + \theta Z_i + \varepsilon_{ij} \quad (19)$$

The dependent variable is the natural log of individual i 's quarterly earnings in employment spell j . The elasticity variable represents the gender specific elasticity of firm j and gender g . 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, education⁶, race, ethnicity, year 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.

6 Results

Summary Statistics

Table 1 reports summary statistics for both men and women in my sample. Since the unit of observation is the employment spell, and only dominant jobs are included, some statistics deviate slightly from typical observational studies of the labor market. The average

⁶Reported 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 model, this is of little concern.

employment spell lasts about two and a half years, with more than sixty percent of spells resulting from a move from another job. Of particular importance to this study, is that the raw earnings gap between men and women is about 0.34 log points. The quarterly nature of the LEHD data make it difficult to precisely identify⁷ 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).

To give the reader some intuition about the type of firms in my sample, the median firm employs roughly 400 workers, hiring 75, in a given quarter. keep in mind that these statistics are not point in time calculations, but rather totals throughout an entire quarter. Additionally, remember that these are at the firm (state-level) rather than at the establishment (individual unit) level.

Firm-Level Measure

Table 2 presents the elasticities estimated through Equations (16)-(18) broken down by gender. The first four columns report the average (weighted by employment) 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 to obtain the labor supply elasticity.

The results detailed in Table 2 are notable in two regards. First, the average labor supply elasticities (0.94 for women and 1.09 for men) are fairly monopsonistic, implying a high degree of market power for firms⁸. This is consistent with previous work utilizing dynamic labor

⁷The 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.

⁸A number of robustness check (equivalent to Webber (2012)) were run to test for threats to identification such as endogenous mobility. No significant differences in the estimated labor supply elasticities were found under any of these alternative specifications. Results are available upon request.

supply models such as Manning (2003) or Webber (2012). It is important to note that Webber (2012) finds that firms do not appear to exploit all of their wage-setting power. Second, the difference between the male and female labor supply elasticities is considerable (1.09 to 0.94)⁹, with the gap implying men should earn approximately 7.5% more than women solely as a result of the disparity in labor supply elasticities¹⁰. This corresponds to about 22% of the raw gender wage gap in my sample, and 33% of the gap when basic observables are controlled for (based on the regressions to be presented below). Finally, we see that the difference in labor supply elasticities between men and women is driven by the difference between the separation and recruitment elasticities to/from employment. In the context of a search model, this implies that the increased search frictions for women are due more to a lower job offer arrival rate as opposed to a higher job destruction rate.

Table 3 displays information about the distribution of labor supply elasticities for men and women in two different samples. The first sample, the same used in Table 2, represents all men and women for whom I was able to estimate a labor supply elasticity (given the restrictions mentioned in the data section). The second sample only includes individuals who work at firms where I am able to estimate both a male and female labor supply elasticity. As shown in Table 3, there is only a small gender differential when looking within firms. Thus nearly the entire elasticity gap between men and women is driven by differences across firms, with women disproportionately working at low-elasticity (and therefore low-wage) firms. This conforms with predictions from the early gender differential literature (Blau, 1977; Groshen, 1991) and the equilibrium search model of Bowlus (1997).

Table 4 reports average labor supply elasticities broken down by NAICS sector. The most competitive industries among men are the manufacturing and mining/oil/natural gas sectors, while the least competitive are the administrative support and accomodation/food service sectors. Among women, the most competitive industries are manufacturing and transportation, while the least competitive are the administrative support and health care

⁹Interestingly, this gap has remained nearly constant throughout the timeframe of my sample.

¹⁰Calculated using Equations (14) and (15)

sectors. The low elasticity for female healthcare workers is consistent with the focus of most of the early monopsony literature’s focus on the market for nurses. The male labor supply elasticity is greater than or equal to the female labor supply elasticity in 18 of the 20 sectors, and only slightly smaller in the other two. By far, the greatest elasticity differential can be found in the construction industry, where men face an elasticity of 1.39 compared to 0.92 for women. Some of these differences are undoubtedly due to differences in occupation within each industry classification, unfortunately the LEHD data do not allow for the identification of occupation.

Now that we have estimated the gender elasticity gap, we now turn to the question of how much of the gender earnings gap can be explained by the difference in labor supply elasticities. Previous studies, which only were able to estimate elasticities at the market level, were forced to interpolate the impact on the gender pay gap. As mentioned above, the theoretical impact implied by my results is men earning 7.5% more than women due to differences in search frictions. Table 5 presents a series of log-earnings regressions¹¹ which allow me to directly estimate this impact due to the firm-level nature of the elasticities generated by this study. In the model with the most detailed set of controls (both person and firm fixed-effects) I find a coefficient of 0.20 on the gender-specific labor supply elasticity, which implies that a labor supply elasticity differential of 0.15 will lead to a gender earnings gap of 3.3%, less than half of the theoretically predicted value¹². This corresponds to about 10% of the raw gender wage gap in my sample and 14.3% of the gender wage gap after controlling for the observables available in this study.

Furthermore, I decompose the labor supply differential by several demographic characteristics. I find that roughly 60% of the gender differential can be explained by marital status and the presence of children. Notably, being married reduces the elasticity of the labor supply curve faced by women by .03 points and having children living at home reduces it

¹¹Table 5 depicts regressions run on the sample of workers who work at firms where both a male and female labor supply elasticity can be estimated. These regressions were also run on the entire sample, as well as on only the male and female samples, with nearly identical results.

¹² $(\exp(.2) - 1) * .15$

a further .06 points. There is no corresponding marriage or child penalty in mobility evident in the male labor supply elasticities¹³.

Also of note in Table 5 is how the coefficient on the gender-specific labor supply elasticity variable changes as person and firm fixed effects are added. The noticeable increase in the coefficient, both when firm and person effects are added to the model, imply that on average low-wage firms have higher labor supply elasticities, and low-wage workers have higher labor supply elasticities. This is in line with the current thinking regarding monopsony power and its interaction with skilled and unskilled labor (Stevens, 1994; Muehleman et al., 2010).

There is reason to believe that the estimates in Table 5 are lower bounds of the true impact of firm monopsony 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. Nevertheless, the measurement error present is unlikely to be of the magnitude necessary to attenuate the estimate by more than half.

While these results are clear evidence that firm-level monopsony contributes to the gender pay gap, as has been documented on a less detailed scale by several other studies, these results provide two key insights into the impact of imperfect competition on the gender pay gap. First, on average, the gap between the male and female labor supply elasticities is quite small within firms which employ a nontrivial number of both men and women. Instead, the gap is primarily driven by disproportionate numbers of men (women) working at high (low) elasticity firms. A second important contribution of this study is that firms do not utilize all of the wage setting power available to them when it comes to the gender pay gap. The

¹³While the LEHD data in general do not allow for identification of marriage or children, the data can be linked with point in time measures from the 2000 decennial shortform Census.

results suggest that women would earn about 7.5 percent less than men, holding all else constant, as a result of increased search frictions. However, I find that these search frictions only cost women about 3.3 percent of their earnings relative to their male counterparts (the analagous statistics for firms which employ nontrivial workers of each gender are 2.0 and 0.9 percent respectively). Given the existence of pay equity laws and substantial social pressure promoting gender equality, this result is not surprising. A similar point was first made by Bronfenbrenner (1956), which argued that firms likely possess substantial wage-setting power but are unlikely to exercise all or most of it.

7 Conclusion

The gender pay gap is one of the most studied topics in modern labor economics. Despite this intense focus, only recently have studies considered the impact that imperfect competition in the labor market may have on the gender pay differential. Furthermore, due to data constraints, the recent empirical studies which find evidence of different degrees of search frictions between men and women are unable to directly estimate the impact of these frictions on the gender pay gap.

This study uses linked employer-employee data to estimate the labor supply elasticity facing the firm, separately by gender, for a comprehensive sample of U.S. firms. Using a dynamic model of labor supply, which identifies the labor supply elasticity to the firm off of job to job transitions, I find evidence of substantial search frictions in the economy, with females facing a higher level of frictions than males. However, the majority of the gender gap in labor supply elasticities is driven by across firm sorting rather than within firm differences, a feature predicted by the Bowlus (1997) equilibrium search model, but which has not been previously documented.

On average, I find that males face a labor supply elasticity 0.15 points higher than females, a differential which leads to 3.3% lower earnings for women (or about 14% of the

adjusted gender earnings gap). However, this is slightly less than half of the theoretically implied impact which the previous literature has been forced to rely upon. Furthermore, I find evidence that women face mobility penalties for marriage and children which men do not face, which explains more than half of the gender labor supply elasticity differential.

References

- 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.
- J. Altonji and R. Blank, *Race and gender in the labor market.*, ser. Handbook of Labor Economics. Elsevier, 1999, vol. 3, pp. 3143–3259.
- G. Becker, *The Economics of Discrimination.* University of Chicago Press, 1971.
- M. Bertrand, *New Perspectives on Gender*, ser. Handbook of Labor Economics. Elsevier, 2011, vol. 4, pp. 1543–1590.
- D. Black, "Discrimination in an equilibrium search model," *Journal of Labor Economics*, vol. 13(2), pp. 309–334, 1995.
- F. Blau, *Equal Pay in the Office.* Lexington Books, 1977.
- F. Blau and L. Kahn, *Women's Work and Wages.* Palgrave Macmillan., 2008.
- H. Bonin, T. Dohmen, A. Falk, D. Huffman, and U. Sunde, "Cross sectional earnings risk and occupational sorting: the role of risk attitudes," *Labour Economics*, vol. 14, pp. 926–937, 2007.

- A. Bowlus, "A search interpretation of male-female wage differentials," *Journal of Labor Economics*, vol. 15, pp. 625–657, 1997.
- 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.
- E. Groshen, "The structure of the female/male wage differential: Is it who you are, what you do, or where you work?" *Journal of Human Resources*, vol. 26(3), pp. 457–472, 1991.
- 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. 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.
- 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.
- J. Madden, *The Economics of Sex Discrimination*. Lexington Books, 1973.
- A. Manning, *Monopsony In Motion*. Princeton University Press, 2003.
- S. Muehlemann, H. Pfeifer, G. Walden, F. Wenzelmann, and S. Wolter, "The financing of apprenticeship training in the light of labour market regulations," *Labour Economics*, vol. 17, pp. 751–774, 2010.
- 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.
- J. Schmieder, "Labor costs and the evolution of new establishments," 2010, working Paper.
- M. Stevens, "A theoretical model of on-the-job training with imperfect competition," *Oxford Economic Papers*, vol. 46, pp. 537–562, 1994.
- D. Sullivan, "Monopsony power in the market for nurses," *Journal of Law and Economics*, vol. 32, pp. S135–S178, 1989.
- D. Webber, "Firm market power and the earnings distribution," 2012, working paper.

Table 1: Summary Statistics

Variable	Men		Women	
	Mean	Std Dev	Mean	Std Dev
Age	38	14	38	14
Tenure (Quarters)	10.2	11.1	10.1	10.83
Log(Quarterly Earnings)	8.68	1	8.34	0.94
White	0.78	0.42	0.76	0.43
Hispanic	0.15	0.36	0.13	0.33
< High School	0.17	0.37	0.13	0.33
High School Degree	0.29	0.46	0.28	0.45
Some College	0.3	0.46	0.34	0.47
College Degree+	0.24	0.43	0.25	0.43

Table 2: Firm-Level Labor Supply Elasticities

Model	ε_R^E	ε_R^N	ε_S^E	ε_S^N	ε
Male Elasticities					
No Controls	.47	.11	-.47	-.62	.96
Full Model	.54	.13	-.54	-.7	1.09
Female Elasticities					
No Controls	.39	.09	-.39	-.62	.83
Full Model	.45	.1	-.45	-.7	.94

The first row of each panel represents estimates from equations (16)-(18) where the only regressor in each model is log earnings. The second row estimates the same equations, and includes age, age-squared, firm size, along with indicator variables for nonwhite, Hispanic, completing a high school diploma, some college, and college degree or greater, and year effects. 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 facing the firm.

Table 3: Distribution of Estimated Firm-Level Labor Supply Elasticities

	Mean	Percentiles				
		10th	25th	50th	75th	90th
All workers						
Men	1.09	0.22	0.45	0.78	1.24	1.94
Women	0.94	0.23	0.43	0.72	1.08	1.58
Only firms with both elasticities						
Men	0.98	0.23	0.44	0.75	1.15	1.69
Women	0.94	0.26	0.46	0.74	1.08	1.54

*Three separate regressions, corresponding to equations (16)-(18), were estimated separately by gender 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 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 4: Mean Labor Supply Elasticity by NAICS Sector and Gender

NAICS Sector	Male Labor Supply Elasticity	Female Labor Supply Elasticity
Agriculture	1.35	1.25
Mining/Oil/Natural Gas	1.51	1.3
Utilities	1.18	1.03
Construction	1.39	0.92
Manufacturing	1.67	1.66
Wholesale Trade	1.38	1.27
Resale Trade	1.01	0.96
Transportation	1.44	1.38
Information	1.11	1.11
Finance and Insurance	1.13	1.2
Real Estate and Rental	0.99	0.94
Professional/Scientific/Technical Services	1.06	1.03
Management of Companies	1.08	1.04
Administrative Support	0.64	0.64
Educational Services	0.95	0.9
Health Care and Social Assistance	0.77	0.75
Arts and Entertainment	0.96	0.87
Accommodation and Food Services	0.76	0.84
Other Services	1.04	0.93
Public Administration	1.28	1.11

*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 (16)-(18), were estimated separately by gender 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 ethnicity, racial status, and education level. Employer controls include the number of employees working at the firm. Year effects are included in all models.

Table 5: Impact of Search Frictions on Earnings

	0.14	0.12	0.08	0.05	0.05	0.06	0.2
Coefficient on labor supply elasticity							
Demographic controls	No	Yes	Yes	Yes	Yes	Yes	Yes
Employer controls	No	No	Yes	Yes	Yes	Yes	Yes
Tenure controls	No	No	No	Yes	Yes	Yes	Yes
State fixed-effects	No	No	No	No	Yes	Yes	Yes
Person fixed-effects	No	No	No	No	No	Yes	Yes
Firm fixed-effects	No	No	No	No	No	No	Yes
R-Squared	0.005	0.233	0.308	0.329	0.336	0.815	0.99

*A pooled national sample of all dominant employment spells, at firms which have estimated elasticities for each gender, 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 are greater than 200 in all models. Clustering these standard errors at various levels does not affect the statistical significance. All standard errors and other estimated coefficients are available upon request. There are approximately 183,000,000 observations in each specification.