I Z A Institute of Labor Economics

Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 12785

Gender Gaps in Pay and Inter-Firm Mobility

Christian Bredemeier

NOVEMBER 2019

DISCUSSION PAPER SERIES

IZA DP No. 12785

Gender Gaps in Pay and Inter-Firm Mobility

Christian Bredemeier
University of Wuppertal and IZA

NOVEMBER 2019

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

Gender Gaps in Pay and Inter-Firm Mobility*

The gender gap in inter-firm mobility is an important contributor to the gender pay gap but is as yet unexplained. In a structural model of workplace choice, I show that the gender mobility gap can be understood as a consequence of women's typical roles as secondary earners in most households which induces households to put more weight on the non-pay dimensions of women's workplaces. I provide direct empirical evidence for this explanation by documenting that the sensitivity of quits to wages is weaker the less an individual contributes to household earnings. Furthermore, gender differences are small once differences in earner roles are accounted for. My quantitative model evaluations show that ignoring the influence of earner roles on inter-firm mobility leads to substantial biases in wage-gap decompositions and predicted policy effects.

JEL Classification: J42, J16, J62, J71

Keywords: labor-market monopsony, gender gaps, job mobility, discrimination

Corresponding author:
Christian Bredemeier
University of Wuppertal
Schumpeter School of Business and Economics
Gaußstraße 20
42119 Wuppertal
Germany

E-mail: bredemeier@wiwi.uni-wuppertal.de

* I wrote most of this paper while I was at the University of Cologne to which I express my gratitude. I further thank Tobias Föll, Falko Jüßen, Michael Krause, Roland Winkler, seminar participants at the universities of Bielefeld, Cologne, Munich (LMU), Witwatersrand (Johannesburg), and Wuppertal and the RWI (Essen) and IZA (Bonn) institutes as well as conference participants at Computation in Economics and Finance, Econometric Society, Economic Association of South Africa, German Economic Association, French Economic Association, and Royal Economic Society for many helpful comments and suggestions.

1 Introduction

The gender pay gap is an important issue for society and an interesting phenomenon for economists seeking to understand the functioning of the labor market (Goldin et al., 2017, Blau and Kahn, 2017, Juhn and McCue, 2017, Goldin, 2014, Beaudry and Lewis, 2014, Bailey et al., 2012, Aizer, 2010). While the gender wage gap has closed substantially over the last 50 years, a substantial gap in wage rates of about 20% still remains in many developed economies. There also is a substantial part of the gender wage gap that cannot be explained by observable characteristics of women and men. Blinder-Oaxaca decompositions of the gender wage gap usually attribute about half of the gender gap to unexplained factors (Blau and Kahn, 2017).

An important contributor to the unexplained gender wage gap is the gender gap in inter-firm mobility which allows firms to discriminate monopsonistically against women and pay them less than men. It is well documented for the United States (e.g., Ransom and Oaxaca, 2010, Webber, 2016, Mas and Pallais, 2017, Wiswall and Zafar, 2017) and other developed economies (e.g., Barth and Dale-Olsen, 2009, Hirsch et al., 2010, Booth and Katic, 2011, Sulis, 2011, Redmond and McGuinness, 2019) that the workplace choice of women, i.e., for which firm a woman works, is, on average, less sensitive to wage differences between firms than the workplace choice of men. Put differently, when choosing between two jobs that differ in pay and in non-pay dimensions such as location, flexibility, or tasks, a man is more likely than a woman to choose the better paying job. Conversely, women are less likely attracted to a new firm by better pay and in this sense less mobile between firms. Firms can exploit this lower mobility of women between firms and pay less to women than to men (for evidence, see Ransom and Oaxaca, 2010, and Félix and Portugal, 2017).[1] The potential contribution to the gender wage gap of the gender gap in inter-firm mobility – and the monopsonistic discrimination against women it allows – is quite substantial. In a survey, Hirsch (2016) summarizes the empirical literature by stating that 40-65% of the unexplained gender

[1] In the U.S. and many other developed countries, it is forbidden by law to condition pay on gender. Firms paying less to women nevertheless occurs in various ways. It is documented that women are promoted less frequently (Bosquet et al. 2019; Gobillon et al. 2015; Pekkarinen and Vartiainen 2006; Booth et al. 2003), paid less on a given rank (Blackaby et al. 2005), receive pay raises less often (Babcock and Laschever 2009; Artz et al. 2018), are assigned to lower-paying jobs (Ransom and Oaxaca, 2005) or tasks (Babcock et al., 2017), and are paid less in discretionary compensation components such as bonuses (e.g., Grund, 2015). The gender gap in inter-firm mobility implies that profit-maximizing firms are able to treat women this way without risking the loss of many female workers. Further, the gender gap in inter-firm mobility implies that women leave low-pay firms at lower rates than men such that in the long run men work at better paying firms (Card et al. 2018; Bayard et al. 2003).

gap can be attributed to this form of discrimination. Despite its importance, the gender gap in inter-firm mobility is not understood, ”evidence on the causes of the gender difference in the wage sensitivity of workers’ labor supply to a single employer is still missing” (Hirsch, 2016, page 8).

In this paper, I provide a theoretical explanation for the gender gap in inter-firm mobility, present direct empirical evidence for it, and assess its implications for the gender pay gap in a quantitative general-equilibrium model. In order to understand why women are less mobile between firms, one has to take into account that most workers live in couples.[2] My explanation for the gender gap in inter-firm mobility is based on a joint decision-making process of households. While the literature has increasingly recognized that labor-supply decisions of married individuals should be considered as joint decisions of the family, most papers focus on the choice of whether or how much to work.[3] I expand this notion to the choice of where to work and whether to change jobs.[4] I focus on workers’ weighing of pay and non-pay characteristics of different jobs when choosing between them and, in my model, spouses decide jointly how much weight to put on these two dimensions.

Consider a double-earner household where income potentials of the spouses differ, for example because they differ in labor-market experience which is valued by firms, such that there is a primary and a secondary earner. In such a situation, a certain percentage increase in the secondary earner’s income is less relevant to the household than the same percentage increase in the primary earner’s income. In turn, this makes non-pay characteristics more important relative to pay when deciding about where the secondary earner should work. I formalize these relations in a model of workers’ choices between heterogeneous workplaces that is similar to the models considered by Card et al.

[2] Between 2005 and 2012, 70% of people aged 35-64 in the U.S. were married, spouse present, or cohabitating with a partner as an unmarried couple (Census Table AD-3). Two thirds of married couples between ages 35-55 were double-earner households.

[3] There is strong empirical evidence (e.g., Cherchye et al., 2012, Donni and Moreau, 2007) that labor supply is a joint decision of spouses and understanding them as such has been helpful to better understand phenomena such as consumption insurance against wage-rate shocks (e.g., Blundell et al., 2016, Autor et al., 2019) and the distribution of female labor supply (e.g., Bick, 2016, Bredemeier and Jüßen, 2013) or normative issues such as optimal unemployment insurance (Ortigueira and Siassi, 2013) or pension systems (Nishiyama, 2019).

[4] The joint-search literature studies the joint decisions of households whether a spouse searches for a job and when to accept a job offer. As is common in the search-and-matching framework, most papers (e.g., Mankart and Oikonoumou, 2017, Wang, 2019) in this literature model heterogeneity of workers and jobs and the matching between the two as embedded in the black-box matching function. Hence, a job and a worker can fit each other (a match) or not and this is determined exogenously while, in my model, I study the decision of households how to balance pay and non-pay characteristics of a job. Guler et al. (2012) have an explicit non-pay dimension of jobs, their location. In this dimension, couple households have a clear incentive to search for jobs for husband and wife which are at similar locations. Importantly, the costs of different job locations are borne equally by both spouses. I understand non-pay job characteristics more broadly and take seriously the issue that a household may pick jobs with differently likeable job characteristics for husband and wife.

(2018) and Wiswall and Zafar (2017) and extends them by introducing double-earner couples. I derive the wage sensitivity of a worker's employer choice as a function of the share that this worker contributes to household earnings. This way, one can interpret the gender gap in inter-firm mobility as reflecting the gap between primary and secondary earners. This is due to the fact that, statistically, women are contributing lower average shares to household earnings than men.

My analysis suggests that pay is a less important determinant of the choice of women's workplaces not primarily because women care intrinsically more about non-pay characteristics of a job but because the household does not rely so much on their earnings.[5] In line with my explanation based on within-household earner roles, Webber (2016) documents that the gap in inter-firm mobility is larger between married men and married women than between singles of both genders, for whom my channel is absent. A similarly supportive finding is reported by Ransom and Oaxaca (2010) who document that the gender gap in inter-firm mobility has become smaller over time, as has the gender earnings gap. Both observations are in line with my channel while the alternative explanation of deep intrinsic gender differences in the importance of non-pay job characteristics would imply that the gender gap in inter-firm mobility should be expected to be rather constant across the population and over time.

For a formal empirical test of my channel, I derive a testable prediction. Specifically, my model predicts that the wage sensitivity of quits is stronger for people who contribute larger shares to household earnings. I test this prediction by estimating quit regressions where I include the share of contributed earnings and interact it with the wage rate.[6] The empirical results are strongly supportive of my theory. I find a significant, quantitatively important, and robust effect of intra-household earner roles on the wage sensitivity of quits in the expected direction. My empirical results imply that an increase of an individual's contributed share to household earnings by ten percentage points raises the individual's mobility between firms by about 10 percent. Once I take

[5] Hirsch et al. (2010) have brought forward verbally a similar reasoning arguing that domestic responsibilities may prevent women from following monetary considerations when choosing their employers and that they may instead care more about the job's location or the working hours offered. While this is clearly a related thought because domestic responsibilities and earner roles are closely related, Hirsch et al. (2010) neither model their argument formally nor investigate it empirically, both of which I do.

[6] Going back to Manning (2003), quit regressions are applied to measure inter-firm mobility as the responsiveness of quits to wage rates. The literature has also considered the share of hires from employment to allow for potential asymmetries between inflow and outflow margins when estimating the elasticity of labor supply to individual firms (Hirsch et al. 2010; Hirsch and Jahn 2015; Hirsch et al. 2018, for example). I restrict attention to the quit margin (as, e.g., Ransom and Oaxaca, 2010, and Depew and Sørensen, 2013) because this allows me to include detailed information on a worker's family background in the analysis.

earner roles into account, the remaining gender differences in the wage sensitivity of quits are small and statistically insignificant. This implies that men and women would be very similar in terms of inter-firm mobility if they contributed equal shares to household earnings.

Regarding the consequences for the gender wage gap, my results imply a mutually enforcing cycle between gender gaps in pay and inter-firm mobility. While I find that the gender gap in inter-firm mobility is in large part a consequence of the gender earnings gap, it is a well-known result that the gender gap in inter-firm mobility is itself a contributor to the gender wage gap. Taken together, this constitutes a feedback mechanism which amplifies any exogenous differences between genders. For example, if an experience gap between men and women makes firms willing to pay higher wage rates to men than to women, this initial wage gap is amplified because households will respond to it by supplying female labor less elastically to individual firms which in turn allows firms to reduce women's wages further. Similar mechanisms amplify initial gender differences in labor supply or non-pay job preferences. I provide a quantitative model assessment that shows that the cycle between gender gaps in pay and inter-firm mobility amplifies the effects on the gender wage gap of changes in exogenous gender differences or gender-specific policy changes by about 30%.

As a consequence of this amplification, a model that takes into account this feedback between the gender gap needs smaller exogenous differences between men and women in order to generate gender gaps in outcomes and, in turn, explains larger fractions of these gender gaps endogenously through economic mechanisms. By contrast, models that mistake inter-firm mobility and gender differences therein as exogenous yield biased estimates concerning the importance of different contributors to wage gaps. My quantitative analysis indicates that taking into account the effect of household earner roles for inter-firm mobility reduces the gender differences in non-pay job preferences required to match observed gender gaps by a factor of about 5 and increases the share of the gender wage gap that should be assigned to labor-demand factors such as experience by about one third.

The remainder of this paper is organized as follows. Section 2 derives the main theoretical results analytically. Section 3 presents the empirical analysis. Section 4 discusses the quantitative model evaluations. Section 5 concludes.

2 Basic model

I first consider a simple model with the minimal set of ingredients to demonstrate the basic mechanism behind my results. This basic model extends the model used by Card et al. (2018) by introducing couple households.[7]

In order to keep the basic model as simple and transparent as possible, I start with a static set-up that includes only the monopsonistic friction. A dynamic monopsony representation that delivers identical steady-state results can be found in Appendix A.1. I address search costs in Appendix A.3. The basic model also abstracts from further aspects of labor supply, job choices, and labor-market competition (such as endogenous hours choices, segregated labor markets, within-gender inequality, singles, gender differences in elasticities of labor supply to the market, home production, and firm entry) which I will address in the quantitative model analysis in Section 4.

The basic model can be described as follows. There is a finite number V of firms and two types of workers. Firms differ in non-pay characteristics over which workers have heterogeneous preferences. This gives firms local monopsony power over those workers who like the respective firm's characteristics. Firms cannot observe an individual worker's job preferences which rules out wage discrimination against individual workers. Each firm posts a pair of group-specific wage rates which workers can costlessly observe. Based on these wage rates, households choose a firm for each worker and firms hire any worker who is willing to accept a job at the offered wage rate. Unlike Card et al. (2018), I differentiate between male and female workers and, decisively, I assume that they live together in couple households and take workplace choices jointly.[8]

The key deviation from the perfect-competition benchmark lies in the importance of non-pay job characteristics. There are many non-pay attributes that affect the attractiveness of a job, including its location, the tasks to be performed, the work climate, the flexibility or predictability of a job, and amenities like child care in the work place. Manning (2011) argues that these non-pay dimensions of a job are the key obstacles to finding a suitable employer and are hence key for understanding what is meant by search on the labor market. Sullivan and To (2014) have shown empirically that such non-pay characteristics play a major role for workers' job search behavior.[9]

[7] A similar modelling of the labor-supply side has been used by Wiswall and Zafar (2017) who also did not consider couple households.

[8] Card et al. (2018) consider workers of different skill groups who live in bachelor households in their model.

[9] While, at first glance, it may seem that many non-pay job characteristics are either good or bad, the model considers

The model allows for potential gender differences in three exogenous dimensions, preference differences regarding the importance of non-pay job characteristics, differences in the marginal revenue product of a worker, and differences in hours worked. While genders have to differ in at least one of these dimensions to generate gender gaps in labor-market outcomes, no exogenous gap is absolutely necessary for the results. For example, the model generates gender gaps in wage rates and inter-firm mobility even if the only exogenous difference between men and women is in their labor supply.

2.1 Households

There are two members in each household, a woman indexed by f and a man indexed by m. They choose jointly how much to consume and where to work. Formally, the household chooses consumption c, and workplaces k_f, k_m for both spouses in order to maximize

$$u = \ln c + \frac{1}{\gamma_f} \cdot \varepsilon_f \left(k_f\right) + \frac{1}{\gamma_m} \cdot \varepsilon_m \left(k_m\right), \tag{1}$$

where the $\varepsilon_g \left(k_g\right)$, with $g = f, m$ indexing gender, describe utility from non-pay characteristics of the chosen firm. I assume that the additional utility agents achieve at the different potential employers, $\varepsilon_g(v)$, are independent draws from a type-I extreme value distribution with scale parameter $1/V$. The exogenous utility weights γ_m and γ_f measure the (inverse) importance of non-pay job characteristics to men and women, respectively. I allow these weights to be gender-specific but the qualitative results will not hinge on this. In fact, the quantitative evaluations show that the weights γ_f and γ_m are rather similar and that most of the gender gap in inter-firm mobility arises endogenously in the model. I scale the distribution of taste shifters ε by the number of firms, reflecting that, in markets where more firms are active, differences between any two firms are smaller and thus can be expected to matter less. This scaling is innocuous in the basic model considered here but affects the interpretation of the quantitative model (see Section 4) where I consider different labor-market segments which differ in V, and affects results in an extension where I consider firm entry which makes V endogenous.[10]

them in a completely neutral way. It takes into account that firms differ in their mix of characteristics and that workers value these mixes differently, in terms of what they like and how important various factors are for them. For example, some people may enjoy a competitive environment while others prefer working in a team.

[10] I abstract from complementarities between the non-pay characteristics of spouses' jobs in utility. An easy way to introduce such complementarities were to let the weight γ_g depend on the realized non-pay job utility of the partner.

Households act subject to the budget constraint

$$c = w_f\left(k_f\right)h_f + w_m\left(k_m\right)h_m,$$

where $w_g\left(k_g\right)$ is the wage rate offered to workers of gender g by the chosen firm k_g and h_g are hours worked by household member g. Hours are exogenous in the basic model but will be endogenous in the quantitative model in Section 4.

Workplace choice. Using standard logit solution techniques (McFadden, 1973), one can easily determine the share of workers working for a firm j among those workers whose spouses work for another firm i. For each gender g, this share is given by

$$\frac{\exp\left(\gamma_g V \ln\left(w_{j,g}h_g + w_{i,-g}h_{-g}\right)\right)}{\sum_{p=1}^{V}\exp\left(\gamma_g V \ln\left(w_{p,g}h_g + w_{i,-g}h_{-g}\right)\right)}, \tag{2}$$

where $-g$ denotes the other gender. Consequently, the total mass of workers of gender g working for firm j is[11]

$$n_{j,g} = \sum_{i=1}^{V} n_{i,-g}\frac{\exp\left(\gamma_g V \ln\left(w_{j,g}h_g + w_{i,-g}h_{-g}\right)\right)}{\sum_{p=1}^{V}\exp\left(\gamma_g V \ln\left(w_{p,g}h_g + w_{i,-g}h_{-g}\right)\right)} \tag{3}$$

As Card et al. (2018), I assume that strategic interaction between firms in wage setting is negligible and consider a symmetric equilibrium.[12] This implies that the slope of the labor-supply curve faced by an individual firm is

$$\begin{aligned}\frac{\partial n_{j,g}}{\partial w_{j,g}} &= \sum_{i=1}^{V} n_{i,-g}\frac{\exp\left(\gamma_g V \ln\left(w_{j,g}h_g + w_{-g}h_{-g}\right)\right)}{\sum_{p=1}^{V}\exp\left(\gamma_g V \ln\left(w_{p,g}h_g + w_{-g}h_{-g}\right)\right)} \cdot \gamma_g \cdot V \cdot h_g \cdot \frac{1}{w_{j,g}h_g + w_{-g}h_{-g}} \\ &= n_{j,g} \cdot \gamma_g \cdot V \cdot h_g \cdot \frac{1}{w_{j,g}h_{j,g} + w_{-g}h_{-g}}.\end{aligned}$$

It follows that the elasticity of labor supplied by workers of group g to an individual firm is

$$\eta_g = \frac{\partial n_{j,g}}{\partial w_{j,g}} \cdot \frac{w_{j,g}}{n_{j,g}} = \gamma_g \cdot V \cdot \frac{w_g \cdot h_g}{w_g h_g + w_{-g}h_{-g}}.$$

However, the realized non-pay utilities are a constant in the symmetric equilibrium. Hence, such an extension would only affect worker's responses to unilateral (off-equilibrium) wage changes by firms in couples where both spouses work for the same firm. Following the literature, I assume that the number of firms is large which makes strategic interaction negligible but also implies that the share of such couples is minimal.

[11] To obtain the total number of workers of gender g at firm i, (2) is multiplied with the number of workers of the other gender at firm i and then summed up over all firms i.

[12] Note that this assumption affects men and women symmetrically.

This elasticity is the key measure of inter-firm mobility as it determines how strongly firms can pay workers below their marginal revenue product, see Section 2.2 below.[13] In Appendix A.3, I show that accounting for search costs would alter the result above only through a uniform re-scaling of the parameters γ_m and γ_f.

Intuitively, a worker who intrinsically does not put much weight on non-pay job attributes (large γ) has a high sensitivity of job choices to pay. Inter-firm mobility further depends on how many firms there are because the number of firms V determines how similar each firm is to the most similar firm in terms of job characteristics. Finally and most importantly, inter-firm mobility depends on the contributed share to household earnings, $w_g h_g/(w_g h_g + w_{-g} h_{-g})$. Earnings of an individual who is married to a partner with very high earnings are of little importance to the household. In the limit, this individual would simply work for the firm that he or she likes best and the reaction to this (or another) firm's pay would be minimal.

The gender gap in inter-firm mobility in my model is

$$\Delta\eta = \ln \eta_m - \ln \eta_f = \Delta\gamma + \Delta w + \Delta h. \tag{4}$$

It hence depends on (deep) gender differences in importance of job characteristics and the gender gaps in hourly wage rates and hours worked. The gender gap in inter-firm mobility this reflects deep gender differences in, say, psychological aspects determining the importance of job characteristics but also endogenous differences stemming from roles in the household. An example for the former can be related to the strong evidence that there are indeed deep gender differences in risk aversion (Croson and Gneezy 2009) which can lead women to dislike competitive work environments more strongly than men (Heinz et al. 2016). An example for the latter type may be working night shifts which presumably men and women both dislike but households may choose to accept this unlikeable job attribute in exchange for higher pay when determining the workplace for the primary earner but not when determining the one for the secondary earner. As I will discuss in the following sections, my empirical and quantitative results suggest that gender differences in inter-firm mobility are driven rather by the latter than the former aspect.[14]

[13] Appendix A.2 presents an alternative derivation of the elasticity of labor supply to individual firms that focuses on worker's quit decisions in response to potential wage cuts by their employers.

[14] Supportively and in the context of the examples above, Manning and Saidi (2010) provide evidence that the gender differences in *preferences* regarding competitive work environment contribute to the gender wage gap in the UK to

Figure 1: Model-predicted elasticity of labor supply to individual firms as a function of the husband's contributed share to household earnings.

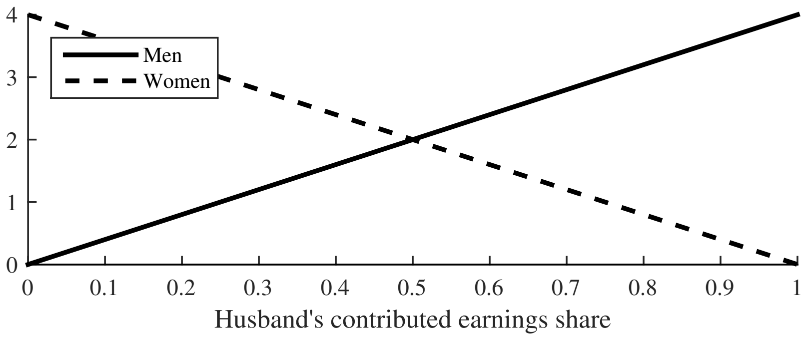

Notes: Example for $\gamma_f = \gamma_m = 0.4$ and $V = 10$.

Figure 1 illustrates the elasticity of labor supply to individual firms as a function of the husband's contributed share to household earnings. With a growing contribution of the husband, the husband's earnings become more important for the household and the wife's earnings become less important. Therefore, the husband's inter-firm mobility increases and the wife's decreases in the husband's contribution. For the figure, I assume that there are no gender differences in the preference parameter γ, hence men and women intrinsically care the same about job characteristics. Consequently, there are no deep gender differences in terms of inter-firm mobility which is reflected in the symmetry of the two lines. If men and women contributed equally to household earnings, the model would predict them to be equally mobile between firms. However, in situations where their earnings differ, the model predicts the primary earner to be more mobile between firms than the secondary earner. Empirically, most couples are in the right part of the figure where men earn more and are more mobile between firms. However, the model predicts that, if the situation were reversed with women earning more than men, women would be more mobile between firms.

2.2 Firms

Firms produce output with labor of both genders. In the production process, a firm cares about total hours work by gender at this firm, the product of the mass of workers attracted by the firm, $n_{j,g}$, and hours worked per worker, h_g. Firms choose wage offers to women and men to maximize

$$a_f n_{j,f} h_f + a_m n_{j,m} h_m - w_{j,f} n_{j,f} h_f - w_{j,m} n_{j,m} h_m$$

a very limited degree. McGee et al. (2015) report similar evidence for the United States.

where a_g is the marginal revenue product of workers of gender g, subject to the labor-supply schedule they face from workers of both genders (described by (3) for $g = f, m$).

As is well known in models of monopsonistic labor markets, it is profit-maximizing for firms to offer a given group wage rates at a mark-down below that group's marginal revenue product with the mark-down depending on the group's elasticity of labor supply to the individual firm. The first-order condition for $w_{j,g}$ is $a_g h_g \cdot \partial n_{j,g}/\partial w_{j,g} - n_{j,g} h_g - w_{j,g} h_g \cdot \partial n_{j,g}/\partial w_{j,g} = 0$ and implies

$$w_{j,g} = a_g \cdot \frac{\eta_g}{1 + \eta_g}. \tag{5}$$

From this behavior of firms, it follows that the gender wage gap in a symmetric equilibrium is, up to first order, given by

$$\Delta w = \Delta a + \frac{1}{1 + \overline{\eta}} \Delta \eta \tag{6}$$

where $\overline{\eta}$ is the average of the gender-specific elasticities of labor supply to individual firms, η_m and η_f and the last step uses a first order Taylor approximation around $\eta_m = \eta_f = \overline{\eta}$. Hence, the gender wage gap is a combination of gender differences in marginal revenue products and monopsonistic discrimination against women.

Note that a gender gap in marginal revenue products, Δa, does not necessarily mean that women per se are worse in performing certain tasks. It should be understood as reflecting all factors outside of the model that lead to a lower contribution of female workers to firm revenue such as differences in physical strength but also the assignment of men and women to different tasks (Babcock et al. 2017) or their segregation on industries with different market power of firms on the goods market. From a human capital perspective, Δa may also reflect foregone experience accumulation during career interruptions, which occur more frequently for women, or lower training investment by firms anticipating such interruptions. A certain part of Δa can also be interpreted as taste-based, or Beckerian, discrimination since discriminatory firms may survive on monopsonistic labor markets. With this broad interpretation of Δa in mind, I will refer to it as the gender gap in labor-demand factors.

2.3 Equilibrium gender gaps

I now briefly describe equilibrium gender gaps before presenting a reference model that does not include couple households. The main equilibrium results are derived in Section 2.5 through a comparison of both model versions.

Equilibrium gender gaps are found by solving the system of equations (4) and (6) for the endogenous gaps Δw and $\Delta\eta$:

$$\Delta w = \frac{\overline{\eta}+1}{\overline{\eta}} \cdot \Delta a + \frac{1}{\overline{\eta}} \cdot \Delta\gamma + \frac{1}{\overline{\eta}} \cdot \Delta h \tag{7}$$

and

$$\Delta\eta = \frac{\overline{\eta}+1}{\overline{\eta}} \cdot \Delta a + \frac{\overline{\eta}+1}{\overline{\eta}} \cdot \Delta\gamma + \frac{\overline{\eta}+1}{\overline{\eta}} \cdot \Delta h. \tag{8}$$

The model generates a mutually enforcing cycle between gender gaps in pay and inter-firm mobility. On the one hand, firms' monopsonistic wage setting induces the gender wage gap to be affected by the gender gap in inter-firm mobility, see (6). On the other hand, deriving the wage sensitivity of workplace choices from a household labor-supply problem has shown that the gender gap in inter-firm mobility depends on the gender earnings gap, see (4).

2.4 A reference model without couples

In order to contrast the results of my model to a model that treats inter-firm mobility as exogenous, I now consider a variant of my model without couples. In this model version, every worker lives alone and seeks to maximize

$$u_g = \ln c_g + \frac{1}{\gamma_g} \cdot \varepsilon_g\left(k_g\right)$$

subject to

$$c_g = w_g\left(k_g\right) h_g.$$

It is straightforward to show that

$$\eta_g = \gamma_g \cdot V \cdot \frac{w_g h_g}{c_g} = \gamma_g \cdot V \tag{9}$$

in this model version. It follows that the gender gap in inter-firm mobility is $\Delta\eta = \Delta\gamma$ in this model. Hence, the gap in inter-firm mobility is solely determined by exogenous preference differences between men and women in this model version and, in this sense, exogenous.

Given group-specific values of the elasticity of labor supply to individual firms, firms' behavior is the same as in the full model. Accordingly, wage offers and the gender wage gap are described by equations (5) and (6) as well. It follows that the equilibrium gender gaps in pay and inter-firm mobility are given by

$$\Delta w = \Delta a + \frac{\overline{\gamma}}{2} \cdot \Delta\gamma, \tag{10}$$

$$\Delta\eta = \Delta\gamma \tag{11}$$

in this model version. Also in this model, the gender wage gap is a combination of gender differences in marginal revenue products and monopsonistic discrimination against women reflecting the gender gap in inter-firm mobility. The latter, however, is here only determined by exogenous gender differences in preferences.

2.5 Main results

I obtain the main qualitative results of the model analysis from a comparison of the full model with couples and the reference model without couples. Here, I obtain qualitative results analytically while, in Section 4, I quantify these results in a larger model that I solve numerically.

Role of preference differences for the gender gap in inter-firm mobility. Comparing the results for the gender gap in inter-firm mobility in both model versions, (4) and (11), the importance assigned to preference differences between men and women becomes apparent. In order to generate a given observable gender gap in inter-firm mobility, $\Delta\eta$, the reference model requires a stronger gender gap in the importance of non-pay job characteristics, $\Delta\gamma$, than the full model that takes into account the joint decision making of double-earner couples. The full model with couples generates a gender gap in inter-firm mobility even without any gender differences in preferences, $\Delta\gamma = 0$, while the reference model strictly requires such differences, $\Delta\gamma > 0$. As a preview of the results of Section 4, the quantitative analysis will reveal that gender differences in preferences are actually quite small.

Amplification. Endogenizing inter-firm mobility through modelling the joint workplace choice of double-earner couples leads to an amplification of changes in exogenous gender gaps. For con-

Table 1: Equilibrium gender gaps in pay and inter-firm mobility in the model versions with and without couples

	full model (endogenous mobility)	**reference model** (exogenous mobility)
gender gap in...		
... pay, Δw	$\frac{\eta+1}{\eta} \cdot \Delta a + \frac{1}{\eta} \cdot \Delta\gamma + \frac{1}{\eta} \cdot \Delta h$	$\Delta a + \frac{1}{1+\eta} \cdot \Delta\gamma$
... inter-firm mobility, $\Delta\eta$	$\frac{\eta+1}{\eta} \cdot \Delta a + \frac{\eta+1}{\eta} \cdot \Delta\gamma + \frac{\eta+1}{\eta} \cdot \Delta h$	$\Delta\gamma$

venience, Table 1 repeats the equilibrium gender gaps in both model variants. Comparing how the exogenous gaps affect the endogenous gaps in both models reveals the amplification in the full model: The effect of a change in any exogenous gender gap, Δa, $\Delta\gamma$, or Δh, on the endogenous gender gaps, Δw and $\Delta\eta$, is stronger in the full model than in the reference model.

The intuition behind this result lies in the mutually enforcing cycle between gender gaps in pay and inter-firm mobility in the full model. For example, a reduction in gender differences in labor-demand factors decreases the gender wage gap directly but through this also increases the importance of women's earnings for households. This makes the choice of women's workplaces more sensitive to pay, i.e., women more mobile between firms. This reduces firms' ability to suppress women's wage rates below marginal revenue product and thus exerts a second, indirect effect on the wage gap. This then again affects the mobility gap and so on.

Similarly, a reduction in the importance of non-pay characteristics of women's workplaces makes women more mobile between firms directly. But the resulting loss in firms' market power vis-à-vis women reduces the wage gap which then again influences the mobility gap and so on.

These amplification mechanisms imply biases in the results of conventional models of monopsonistic labor markets. Overlooking the endogeneity in inter-firm mobility that results from the joint workplace choices in double-earner couples leads to an underestimation of the effects that changes in exogenous gender differences have on gender gaps in endogenous labor-market outcomes.

Decomposition of gender gaps. My model analysis suggests an additional driver of pay differences between men and women. Unlike in the reference model, in the full model there is a role of labor-supply differences between men and women for the gender wage gap. When women supply less labor than men, their earnings are less important to households than men's. It follows that non-pay job characteristics play a larger role in the choice of women's workplaces and women

Table 2: Exogenous gender gaps required to match observable gender gaps, basic model.

	full model (endogenous mobility)	**reference model** (exogenous mobility)
Δa	$\Delta w - \frac{1}{1+\eta} \cdot \Delta\eta$	$\Delta w - \frac{1}{1+\eta} \cdot \Delta\eta$
$\Delta\gamma$	$\Delta\eta - \Delta w - \Delta h$	$\Delta\eta$

become less mobile between firms than men. Monopsonistically competitive firms exploit this and pay lower wage rates to women. Hence, in my model, a part of the gender wage gap is due to households supplying less female than male labor which in turn may reflect absolute advantages of women in non-market activities such as child-rearing or breast-feeding or social norms regarding gender roles in the family. A model that neglects the endogeneity of inter-firm mobility which arises from joint decision making of double-earner households overlooks this channel.

In order to study the importance of the different exogenous gender gaps for the endogenous gaps in pay and inter-firm mobility, I perform a calibration exercise and then decompose the gender gaps within the calibrated model versions. In the basic model, this allows to understand the intuition behind important results. In Section 4, I repeat this exercise within the quantitative model.

The two competing model versions match observable values for Δw and $\Delta\eta$ for the exogenous gender gaps summarized in Table 2. While the calibrated value for Δa is the same for both versions, the full model requires smaller gender difference in preferences for job characteristics than does the reference model (as long as there is a gender gap in earnings, $\Delta w + \Delta h > 0$), as discussed above.

I will now analyze how important the three dimensions of exogenous differences between men and women are for explaining the endogenous gender gaps in pay and inter-firm mobility. Using the calibration results in the equations describing the equilibrium gender gaps Δw and Δh, I obtain the percentage contributions of the different exogenous gender gaps for the endogenous gaps. For example, I calculate the contribution of labor-demand differences for the gender gap, $\Delta w|_{\Delta a}/\Delta w$, by determining the gender wage gap when the other two exogenous gender gaps are counterfactually set to zero. The results of this decomposition exercise are summarized in Table 3.

The results above show that endogenizing inter-firm mobility through incorporating couple households in the model increases the importance of labor-demand factors and reduces the im-

Table 3: Decomposition of endogenous gender gaps in the basic model.

	full model (endogenous mobility)	**reference model** (exogenous mobility)
Share of wage gap explained by...		
... labor-demand differences, $\Delta w\vert_{\Delta a}/\Delta w$	$\frac{\eta+1}{\eta}\cdot\frac{\Delta w-\frac{1}{1+\eta}\Delta\eta}{\Delta\eta}$	$\frac{\Delta w-\frac{1}{1+\eta}\Delta\eta}{\Delta\eta}$
... preference differences, $\Delta w\vert_{\Delta\gamma}/\Delta w$	$\frac{1}{1+\eta}\cdot\frac{\Delta\eta-\Delta w-\Delta h}{\Delta w}$	$\frac{1}{1+\eta}\cdot\frac{\Delta\eta}{\Delta w}$
... labor-supply differences, $\Delta w\vert_{\Delta h}/\Delta w$	$\frac{1}{\eta}\cdot\frac{\Delta h}{\Delta w}$	0
Share of mobility gap explained by...		
... labor-demand differences, $\Delta\eta\vert_{\Delta a}/\Delta\eta$	$\frac{\eta+1}{\eta}\cdot\frac{\Delta w-\frac{1}{1+\eta}\Delta\eta}{\Delta\eta}$	0
... preference differences, $\Delta\eta\vert_{\Delta\gamma}/\Delta\eta$	$\frac{\Delta\eta-\Delta w-\Delta h}{\Delta\eta}$	1
... labor-supply differences, $\Delta\eta\vert_{\Delta h}/\Delta\eta$	$\frac{\eta+1}{\eta}\cdot\frac{\Delta h}{\Delta\eta}$	0

portance of preference differences for understanding gender gaps in pay and inter-firm mobility.[15] Put differently, the result implies that a model that neglects the endogeneity of inter-firm mobility that arises from joint decision making of double-earner households delivers biased estimates of the importance of the exogenous gender gaps in labor demand and preferences. Such a model overestimates the importance of preference differences between genders and underestimates the importance of labor-demand factors.

To clarify, my full model does not yield larger gender differences in labor-demand factors such as productivity or experience. In fact, the calibration of both model versions delivers the same Δa. Rather, the full model explains a larger share of the gender gaps in pay and inter-firm mobility with this given difference between men and women due to the amplifying cycle between the gender gaps in pay and inter-firm mobility discussed above.

2.6 Deriving a testable prediction

I now derive a testable prediction of my theory that I will then assess empirically in Section 3. In particular, I consider quit regressions (Manning, 2003), i.e., regressions of the probability of a worker quitting his or her employer on the worker's log wage rate and appropriate controls. On a monopsonistic labor market, workers leave better paying firms at lower rates and the slope of

[15] In line with the data, this statement assumes that the gender gaps in wage rates, earnings, and inter-firm mobility are all positive, $\Delta w > 0$, $\Delta w + \Delta h > 0$, and $\Delta\eta > 0$. $\Delta w|_{\Delta a}/\Delta w$ is larger in the full model since $(\eta+1)/\eta>1$. $\Delta w|_{\Delta\gamma}/\Delta w$ is smaller in the full model since $\Delta w + \Delta h > 0$. $\Delta\eta|_{\Delta a}/\Delta\eta$ is larger in the full model since it is positive there and zero in the reference model. $\Delta\eta|_{\Delta\gamma}/\Delta\eta$ is smaller in the full model since $\Delta w + \Delta h > 0$.

this relation is closely connected to the elasticity of labor supply to individual firms.[16] My model predicts this slope to depend on household earner roles.

In order to perform quit regressions in my model, I need to incorporate worker flows between firms and wage differences across firms and, to this end, I add two elements to the model, both of which do not alter the results obtained so far. To induce worker flows between firms, I assume that each period a fraction θ of workers is randomly selected and randomly assigned new preferences over firms' non-pay job characteristics. One interpretation of this change in a worker's preferred job characteristics is a change in the family situation such as a child birth, a child moving out of the household, a parent needing care, or the household moving, i.e., changing its geographical location for exogenous reasons such as inheriting a house. To induce wage differences across firms, I assume that firms face small idiosyncratic permanent productivity shocks. Better paying firms employ more workers and are also less likely to be left by workers who experience a change in non-pay job preferences.[17]

In this simple environment, I can derive quit probabilities and the results of quit regressions analytically. Workers quit when their new non-pay job preferences differ sufficiently from their old ones while what is sufficient depends on pay differences between the current and potential new employer. From the perspective of a firm, this means that higher pay not only attracts more workers but also reduces the share of current workers who quit. Formally, the share of workers of group g that leaves firm j between the preceding and the current period is given by

$$q_{g,j} = \theta \left(1 - n_{g,j}\right) \approx \theta \left(1 - \frac{1}{V}\left(1 + \eta_g \ln\left(w_{g,j}/w_g\right)\right)\right) \tag{12}$$

where w_g is the average wage rate paid to workers of group g and the last step uses a first-order

[16] Going back to Manning (2003), quit regressions are applied to measure inter-firm mobility as the responsiveness of quits to wage rates. Some studies also consider the responses of hires from employment to wage rates in order to quantify the recruit elasticity separately from the quit elasticity (Hirsch et al. 2010; Hirsch et al. 2018). Unlike those studies, I concentrate on the quit margin as also done by, e.g., Ransom and Oaxaca (2010), Depew and Sørensen (2013), and Hirsch et al. (2019). The main reason for this decision is that I am going to use detailed information on workers' family backgrounds which I obtain from household survey data. In such data, I can construct information on quits but, by construction, I do not have information on firms' hiring behavior because the unit of observation in the data set is a household, not a firm. By contrast, firm information is available in linked employer-employee data like the German LIAB which is used by, e.g, Hirsch, Jahn, and Schnabel (2018) but this data has no information on family background which is indispensable for my analysis.

[17] In the quantitative analysis performed in Section 4, I consider firm-specific productivity shocks with autocorrelation ρ such that log productivity of gender g at firm j is given by $\ln a_{g,j,t} = (1-\rho)\ln a_g + \rho \ln a_{g,j,t-1} + \xi_{j,t}$, where $\xi_{j,t}$ is the productivity innovation to firm j in period t. With persistent productivity shocks, a firm that pays high wage rates this period also tends to do so in the next period and is therefore less likely to be left. For this reason, future quits are linked to current wage rates.

Figure 2: Model-predicted wage sensitivity of quits as a function of the husband's contributed share to household earnings.

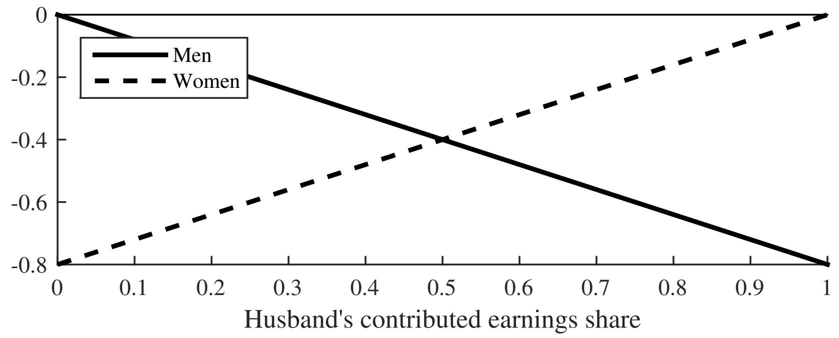

Notes: Marginal effect of $\ln w_{i,g}$ in a regression of (13). Example for $\gamma_f = \gamma_m = 0.4$, $\theta = 0.2$.

Taylor approximation in logs.

I first consider a standard quit regression of the type

$$q_{i,g} = \beta_0 + \beta_1 \cdot \ln w_{i,g} + \gamma \cdot \widetilde{w}_g + \zeta_{i,g},$$

in a random sample of workers of group g indexed by i. q_i describes a subsequent quit by worker i, $w_{i,g}$ is the wage rate currently paid to worker i, $\widetilde{w}_g$ describes the overall (average) pay to workers like worker i, and $\zeta_{i,g}$ is a residual. In such a regression, the estimated coefficient on the log wage rate,

$$\widehat{\beta}_1 = -\frac{\theta}{V}\eta_g,$$

is the larger in absolute value the more mobile the considered group of workers is between firms, i.e., the higher η_g is.[18]

I now extend quit regressions in order to test the main prediction of my model that a worker's inter-firm mobility is determined by the worker's earner role within the household. Recall that my model predicts $\eta_i = \gamma_i \cdot e_i$ where e_i is the worker's contributed share to household earnings, $e_i = w_i h_i / (w_i h_i + w_{-i} h_{-i})$. With this result and (12), it is apparent that a regression of the type

$$q_{i,g} = \beta_0 + \beta_1 \cdot \ln w_{i,g} + \beta_2 \cdot e_{i,g} \cdot \ln w_{i,g} + \beta_3 \cdot e_{i,g} + \gamma \cdot \widetilde{w}_g + \zeta_{ig}, \tag{13}$$

[18]This result implies that the elasticity of the quit rate to the wage rate is $-1/\left(V-1\right) \cdot \eta$. As in all models of monopsonistic labor markets, the elasticity of labor supply to individual firms can be determined as the sum of the (absolute) elasticities of the quit rate and the number of recruits to the wage rate. The latter is $\left(V-2\right)/\left(V-1\right) \cdot \eta$ in my model.

yields a negative coefficient on the interaction term,

$$\widehat{\beta}_2 = -\theta\gamma_i < 0. \tag{14}$$

That is, we should observe a stronger wage sensitivity of quits for workers who contribute larger shares to household earnings, i.e., for whom $e_{it} = w_i h_i/(w_i h_i + w_{-i} h_{-i})$ takes greater values. Intuitively, quits are determined through weighing differences between one's old and potential new job in both pay and non-pay characteristics. Households put much weight on pay when making quit decisions for a person that contributes large shares to household earnings and whose earnings are thus important to the household. This prediction is visualized in Figure 2 where I have chosen an example without gender differences in preference, i.e., $\gamma_f = \gamma_m$, in order to highlight the model's property to generate endogenous gender differences in the wage sensitivity of quits when men's and women's contributions to household earnings differ. The negative relation between the contributed share to household earnings and the wage sensitivity of quits, as summarized in (14), is the key testable prediction of my model.

3 Empirical analysis

In this section, I test the key prediction of my model that inter-firm mobility depends on earner roles in the household. In particular, I test the prediction that the coefficient on the interaction between the log wage rate and the individual's contributed share to household earnings, β_2 in (13), is negative.

3.1 Specification and data

In order to implement the procedure outlined in Section 2.6 in real-world data, I have to address the challenge that alternative wage offers $\widetilde{w}$ are usually not observable. I follow Manning (2003) and the subsequent literature (e.g., Ransom and Oaxaca, 2010, Depew and Sørensen, 2013) and account for alternative wage offers by including a vector X of characteristics that determine a worker's earnings capacity, i.e., the individual's Mincer variables.

Specifically, I estimate

$$q_{ijt} = \Phi\left(\beta_0 + (\beta_1 + \beta_2 \cdot e_{it}) \cdot \ln w_{ijt} + X_{it}\varphi_1 + Y_{it}\varphi_2\right) + \zeta_{ijt},$$

where i indexes the individual, j the firm, and t time, Φ is a monotone $\mathbb{R} \to (0, 1)$ function which can be the identity, logit, or probit function for example. The vector X_{it} collects determinants of the worker's alternative wage offers from other firms. Y_{it} collects terms required to estimate the coefficient on the interaction term consistently such as the non-interacted earnings contribution. ζ_{ijt} is a residual. In Section 2.6, I derived as the main testable prediction of my model that β_2 is negative, implying that the wage sensitivity is larger for greater values of e_{it}. While I start with treating e_{it} as exogenous, I later address the point that e_{it} is arguably endogenous to factors determining earnings potentials including monopsonistic discrimination in IV specifications to isolate exogenous variation in e_{it}.

I use data from the Panel Study of Income Dynamics (PSID).[19] Its panel structure allows me to construct quits because I can observe if a job continues to the next wave and, if it does not, who ended the job, the worker or the firm. Further, since the PSID is a household survey, it contains rich information on workers and their spouses. This information is crucial for my purposes since I want to analyze the impact of household earner roles for worker mobility. However, this information comes at the price of having relatively little information about firms as compared to the case where data from a firm panel or linked employer-employee data is used.

I select a sample of roughly 40,000 jobs ijt held by married individuals between 1978 and 1996 for which I observe wage rates, region, the worker's age, education, race, number of children, and total household earnings, as well as whether the job continued to exist in the next year and, if it does not, the reason for its non-continuation. Requiring the latter information restricts the sample to end in 1996 because the PSID turned biennial in 1997 such that 1996 is the last year for which I know whether jobs still existed the next year. The sample starts in 1978 because the PSID contains information on the reasons of worker-firm separations for both genders continuously from 1979 (hence jobs held in 1978) on. Details on data selection and variable construction can be found in the appendix.

My baseline regression specification is as follows. I estimate gender-specific linear probability models where the dependent variable is a quit of worker i at firm j between years t and $t+1$, q_{ijt}. As main independent variables, I include the log hourly wage rate in year t, $\ln w_{ijt}$ and an interaction

[19] Full data source information: Panel Study of Income Dynamics, public use dataset. Produced and distributed by the Survey Research Center, Institute for Social Research, University of Michigan, Ann Arbor, Michigan, USA.

between this variable and the worker's earnings as a share of total household earnings in this year, $\ln w_{ijt} \cdot (e_{it} - 0.5)$. I measure the earnings contribution as its deviation from earnings parity, 0.5. This choice does not affect the coefficient on the interaction term but allows to interpret the coefficient on the non-interacted log wage rate as the wage sensitivity of quits at earnings parity in the household, i.e., where both partners contribute 50% to household earnings. In order to estimate the coefficient on the interaction term consistently, I also include the contributed share to household earnings as well as a linear-quadratic interaction between the log wage rate and time. The latter interaction is necessary because spouses' contributed shares to household earnings have gender-specific trends reflecting the closure of the gender gap in earnings such that the interaction $\ln w_{ijt} \cdot e_{it}$ could pick up a trend in the wage sensitivity of quits for workers of a given gender rather than the influence of earner roles within the worker's household.[20] In order to control for alternative wage offers, I include in the control vector X year dummies, education dummies, dummies for the number of children, region dummies, and race dummies. This selection of alternative wage controls follows Manning (2003).[21]

I also run specifications where I instrument the earnings contribution e_{it} and account for the potential role of parenthood. In robustness checks summarized below, I estimate the models with non-linear (probit and logit) methods, interact the control variables with the earnings contribution, vary the set of control variables, account for potential non-linear effects of the wage rate, and separately consider the impact of the different variables that enter the construction of the earnings contribution.

3.2 Results

Table 4 summarizes the main results. The table shows the coefficients on the log wage rate and on the interaction term in the baseline specification, separately for men and women. The numbers in brackets are standard errors and asterisks indicate statistically significant difference from zero.

First, I consider a specification without the interaction term between wages and the contributed earnings share, see columns (1) and (2) of Table 4. In line with the literature, I find quite substantial

[20] In a robustness check, I exclude the interaction with time which has only small effects on results.

[21] Following the literature, I estimate regressions separately for men and women as opposed to a joint regression including an interaction of wages with gender because alternative wage offers cannot be expected to depend on the included control variables in the same way for men and women.

Table 4: The wage sensitivity of quits.

	(1) **men**	**(2)** **women**	**(3)** **men**	**(4)** **women**
log wage rate	-0.0488^{***}	-0.0294^{**}	-0.0465^{***}	-0.0469^{***}
	(0.0158)	(0.0164)	(0.0158)	(0.0171)
log wage rate			-0.0389^{***}	-0.0579^{***}
$\times$ (earnings share-0.5)			(0.0130)	(0.0162)
observations	20231	20131	20231	20131
marginal effect of $\ln w$...				
... at mean earnings shares			-0.0531^{***}	-0.0375^{**}
$(e_m = 0.67, e_f = 0.33)$			(0.158)	(0.165)
... at earnings parity			-0.0465^{***}	-0.0469^{***}
$(e_m = e_f = 0.5)$			(0.0158)	(0.0171)

Notes: Dependent variable: quit between observation year and following year; linear probability models; standard errors in parentheses; *, **, *** indicate $p < 0.1$, $p < 0.05$, $p < 0.01$; alternative wage determinants (X, dummies): year, education, region, race, kids; additionally included variables: constant, earnings share, $\ln w \times t$, $\ln w \times t^2$. Variable definitions, sample selection, and specification described in detail in Appendices B.1, B.2, and B.3, respectively.

gender differences in the estimates. The quit behavior of women is, on average, substantially less wage-sensitive than that of men, reflecting their lower inter-firm mobility. Quantitatively, the coefficient for women is about a third smaller than that for men which is similar to the results of Ransom and Oaxaca (2010) and Hirsch et al. (2010).

Columns (3) and (4) report results for my baseline specification including the interaction of wages and the contributed earnings share. The results support strongly the model prediction that inter-firm mobility depends on earner roles within the household. For both men and women, the coefficient on the interaction term is significantly negative which shows that a higher share in household earnings raises the sensitivity of quits to wages as predicted by my model.

Note that the linear effect of the (non-interacted) log wage rate is now very similar across genders. Once different earner roles are accounted for, there remain little gender differences in inter-firm mobility. This indicates that deep gender differences in the importance of non-pay job attributes are likely small. Put differently, my estimation results indicate that women are on average less mobile between firms, not per se but because of their earner roles within the household.

The bottom part of the table gives the marginal effect of the log wage rate at two relevant points in the distribution of contributed earnings shares. At their respective average earnings

Figure 3: Estimated wage sensitivity of quits as a function of the husband's contributed share to household earnings.

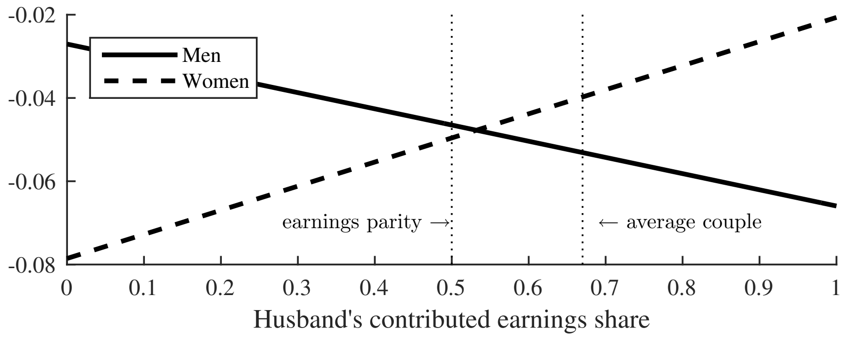

Notes: Marginal effect of $\ln w$.

contributions, women's quitting behavior is substantially less sensitive to wages with numbers being comparable to the linear effect in columns (1) and (2). By contrast, the empirical model predicts that, if husbands and wives contributed equal shares to household earnings, they would also show a similar mobility between firms: at an earnings contribution of 0.5 ("earnings parity"), gender differences in the marginal effect of the wage rate on the quit probability are small.

Figure 3 completes this picture and plots the marginal effect of the log wage rate on the quit probability for different values of the *husband's* earnings contribution (hence, for wives, one minus their own contribution). While the linearity follows from the regression specification, slopes and intersections of the two lines resemble their counterparts from the theoretical model, Figure 2: if their earning shares were similar, men and women would also be similarly mobile between firms. However, since most husbands contribute more to household earnings than their wives, they are more mobile between firms. In line with the theoretical model, the empirical results imply that, if earner roles within households were reversed with men contributing less to household earnings than women, they would be less mobile between firms.

Table 5 contains results of additional estimations. I first account for the fact that household earner roles depend on how much a person could earn, including how strongly this person is discriminated against by monopsonistic firms. Earner roles may endogenously react to mobility between firms and quit probabilities. I use an IV strategy to address this point. Instrumenting a person's earnings share is challenging in the framework of a quit regression because the individual's Mincer variables already have to be included in the second stage to account for alternative wage

Table 5: Accounting for the potential endogeneity of the earnings contribution, the role of children, and transitions out of labor force.

	men	**women**
IV	-0.0457^{*}	-0.0951^{**}
	(0.0320)	(0.0443)
parents only	-0.0279^{*}	-0.0668^{***}
	(0.0160)	(0.0188)
include interaction	-0.0376^{***}	-0.0558^{***}
with no. of children	(0.0004)	(0.0165)
remain in	-0.0338^{**}	-0.0396^{*}
labor force	(0.0172)	(0.0219)

Notes: Results from quit regressions; coefficients on the interaction term between log wage rate and the earnings share; standard errors in parentheses; *, **, *** indicate $p < 0.1$, $p < 0.05$, $p < 0.01$; alternative wage determinants (X, dummies): year, education, region, race, kids; additionally included variables: constant, e, $t \times \ln w$, Xe. Variable definitions, sample selection, and specification described in detail in Appendices B.1, B.2, and B.3, respectively.

offers. However, I can make use of the individuals in my sample living in couples and that the *relative* earnings also depend on the earnings potential of their partners. Specifically, I can use the partner's Mincer variables as instruments for e_{it}. Additionally, I include partner's industry and partner's occupation to instrument the earnings contribution. When I follow this strategy, I find my main results confirmed, see first row if Table 5. The coefficients on the interaction term are significantly negative. Not surprisingly, standard errors are larger as, here, they also include first-stage uncertainty.

Next, I want to rule out that my results are driven by the presence of children in a household. The birth of a child could be an event that leads to a reduction in the mother's earnings through reduced hours worked (and hence a reduction in her contributed share to household earnings) and simultaneously to an increased importance of non-pay characteristics of the mother's job (such as distance to a child-care facility), without the causality running through her earnings share. I address this point in two specifications for which results are shown in the second and third rows of Table 5. First, I restrict the sample to parents with children living in the household. Here, results cannot be driven by differences in earnings shares and importance of non-pay job characteristics between parents and non-parents. In this specification, I still find significantly negative coefficients

Table 6: Robustness checks.

	men	women
probit	-0.0515^{***}	-0.0544^{***}
	(0.0135)	(0.0148)
logit	-0.0544^{***}	-0.0870^{***}
	(0.0135)	(0.0169)
polynomial in $\ln w$	-0.0385^{**}	-0.0833^{***}
	(0.0189)	(0.0190)
interact controls	-0.0383^{***}	-0.0500^{***}
	(0.0134)	(0.0162)
include tenure	-0.0356^{***}	-0.0295^{*}
	(0.0128)	(0.0160)
no time interaction	-0.0449^{***}	-0.0519^{***}
	(0.125)	(0.158)
use 1972 information	-0.354^{***}	—
for men	(0.0126)	—

Notes: Results from quit regressions; coefficients on the interaction term between log wage rate and the earnings share; average marginal effects for probit and logit; standard errors in parentheses; *, **, *** indicate $p < 0.1$, $p < 0.05$, $p < 0.01$; alternative wage determinants (X, dummies): year, education, region, race, kids; additionally included variables: constant, e, $t \times \ln w$, Xe. Variable definitions, sample selection, and specification described in detail in Appendices B.1, B.2, and B.3, respectively.

on the interaction terms. Second, I additionally include an interaction between the number of children and the log wage rate that should pick up the effect of children on the wage sensitivity of job choices independent of relative earnings. Put differently, the coefficient on the interaction of the wage rate and the earnings share measures the effect of earnings on inter-firm mobility conditional on a given number of children. I find my main results confirmed.[22]

One may also be concerned about the role of individuals who quit their job and leave the labor force, for example into retirement or parental leaves. Such quits are included in my baseline empirical analysis but do not occur in my model. For this reason, I also estimate the quit regressions

[22] As one would expect, the coefficient on the interaction between children and the log wage rate is negative for men and positive for women. Hence, children tend to make the choice of fathers' workplaces more wage-sensitive and the choice of mothers' workplaces less wage-sensitive. However, this effect of children is not driving my main results. I also estimated my baseline specification for married couples without children in the household and also find negative coefficients on the interactions of wage rates and earnings shares but estimates are less precise in this relatively small sample that consists mostly of relatively young or relatively old workers.

Table 7: Interaction with alternative variables

	men	**women**
own earnings	-0.0044^{**}	-0.0094^{***}
	(0.0017)	(0.0018)
own wage rate	-0.0047^{***}	-0.0045^{**}
	(0.0017)	(0.0021)
own hours	-0.0031	-0.0108^{***}
	(0.0072)	(0.0028)
partner's earnings	0.0048^{*}	0.0083^{*}
	(0.0024)	(0.0050)
average household earnings	0.0421^{**}	0.1587^{***}
	(0.0178)	(0.0227)

Notes: Results from quit regressions; coefficients on the interaction term between log wage rate and the indicated variable; p-values in parentheses; *, **, *** indicate $p < 0.1$, $p < 0.05$, $p < 0.01$; alternative wage determinants (X, dummies): year, education, region, race, kids; additionally included variables: constant, the indicated variable, $t \times \ln w$. Variable definitions, sample selection, and specification described in detail in Appendices B.1, B.2, and B.3, respectively.

for a sample of individuals who are in the labor force also in the subsequent year ($t + 1$). With these restricted samples, I continue to find significantly negative coefficients on the interaction of the log wage rate and the earnings share, see fourth row of Table 5. This shows that my results indeed reflect the quitting behavior of individuals who either directly transition to a new job or plan to do so in the near future.

To check the robustness of my results, I also consider a number of respecifications of my baseline quit regression, see Table 6. While the baseline approach is a linear probability model, the first and second lines in Table 6 show the marginal effects of the interaction term in probit and logit models, respectively. Also here, the coefficient on the interaction term is significantly negative for both men and women. The third line allows for non-linear effects of the wage rate as in Depew and Sørensen (2013) by including a cubic of the wage rate. This is an interesting extension since the documented effect of the earnings contribution may reflect a non-linear effect of the wage rate. The results show that it does not since the coefficients on the interaction terms are still significantly negative when controlling for higher orders of the wage rate. The fourth line of Table 6 takes into account that workers with higher contributions to household earnings should also react more strongly to alternative wage offers by interacting the alternative wage determinants X with the earnings contribution. Also here, the coefficient on the interaction term remains negative. The fifth line includes tenure as a control variable as discussed by Manning (2003). While this impacts on the precision of the estimates for women, it does not affect substantially the point estimates. The sixth line reports results for a specification where I omit the interaction between the log wage rate and time, which also does not impact on the main results. Finally, the tenth line additionally uses data for 1972, where quits can be constructed for men but not for women. Including these additional data points affects the results for men only marginally.

Since an individual's contributed share to household earnings is a constructed variable, I also perform estimations where I interact the log wage rate with its components. Table 7 reports the results. In line with the theoretical model, workers are found to supply labor more elastically to individual firms (hence, their wage sensitivity of quits, which is negative, is smaller) if, ceteris paribus, they earn more, have higher wage rates, work longer hours, or are married to partners with lower earnings.[23] This corroborates the finding that my main results reflect more than a simple non-linear effect of the wage rate itself. The final line addresses the possibility of intertemporal consumption smoothing. In presence of substantial fluctuations in household earnings, marginal utility of consumption would be determined by average rather than current household earnings if the household has access to sufficient consumption-smoothing possibilities. The results confirm the model prediction that households with higher average earnings supply labor of their members less elastically (the coefficient is positive thereby dampening the negative effect of wage rates on quit probabilities).

4 Quantitative analysis

While the last section has provided empirical evidence for the new link between earner roles and inter-firm mobility I propose, I now quantify my theoretical results regarding the feedback mechanism between gender gaps in pay and inter-firm mobility it implies. To this end, I extend my model to take into account further aspects of employer choices, labor market competition, and the

[23] It is not surprising that the effect of the interaction with hours worked is estimated imprecisely for men due to the low variation of hours worked among men.

gender wage gap. The main extensions are an endogenous hours choice, the inclusion of singles, within-gender heterogeneity, and considering more dimensions of labor supply, i.e., allowing for vertical (e.g., industry or occupation) and horizontal (e.g., education or career) choices of workers. I also address a number of further aspects (such as home production, gender differences in the elasticity of labor supply to the market, firm entry, and the type of wage competition between firms) in additional model versions considered in the Appendix. I calibrate the model and use it for counterfactual experiments such as gender-gap decompositions and comparative statics.

4.1 Additional model features

Vertical and horizontal segregation. Segregation of men and women into different segments of the labor market is an important aspect of the gender wage gap (see, e.g., Blau and Kahn, 2017). Further, this segregation is affected by similar mechanisms as the one regarding workplace choices described in my basic model. E.g., when workers decide in which industry to work, a similar weighing between pay and non-pay characteristics of the industry takes place and, when this is done in the family, the weight that is put on the two dimensions is affected by similar considerations as firm choices considered in the basic model from Section 2.

I incorporate the possibility of vertical and horizontal segregation by allowing workers to choose from horizontal industry-occupation cells that differ in the degree of labor-market competition between firms and from vertical segments of the labor market that differ in a worker's contribution to firms' marginal revenue product. As an example, an individual's horizontal choice could be to concern herself with law rather than finance or medicine. Her vertical choice could be to be a lawyer rather than a paralegal or to invest into becoming a partner rather than remaining at lower levels of the hierarchy. The decision which law firm to work for would be the choice of workplace already considered in the basic model in Section 2.

Formally, (horizontal) differences in the degree of labor market competition are modelled as differences in the scale parameter of the distribution that governs workers' non-pay preferences over firms, $1/V_z$ where z indicates the horizontal segments of the labor market. An interpretation of these differences is that some labor markets are more concentrated such that there are fewer firms and the non-pay differences between any two firms is larger. Further, (vertical) differences in workers' marginal revenue product are achieved by segment-specific multiplicative productivity shifters α_y where y indicates the vertical levels of the labor market. An interpretation is that workers with more schooling work on higher levels of the hierarchy within a firm and are more productive. When calibrating my model, I follow the interpretation of V_z and α_y as firm concentration in an industry and productivity differences between workers with different education, respectively.

Heterogeneity in horizontal and vertical labor-supply choices is achieved through two modelling elements. First, individuals have preferences over the different segments of the labor market like they have preferences over workplaces within these segments. These preferences are also a stand-in for the costs of education as far as they affect which labor-market segment is or can be entered. Specifically, I add further taste shifters, $\epsilon_g(y_g, z_g)$, to utility that reflect the level of utility a worker obtains directly from working in a specific segment of the labor market. Also here, I assume that taste shifters are type-I extreme-value distributed and denote the scale parameter by σ. Second, I allow for within-gender (and within-education) heterogeneity in productivity which affects a worker's marginal revenue product in a firm, i.e., the marginal revenue product of a worker is determined by the individual-specific a_i and the segment-specific α_z and given by $\alpha_z \cdot a_i$.

Endogenous hours choice. While hours worked are exogenous in the basic model, I now incorporate an endogenous choice of labor supply at the intensive margin. I do so because hours worked are an additional margin at which households can react to pay differences between men and women and a factor that determines the importance of hourly pay for workplace choices.

Formally, instead of (1), household preferences are now described by

$$\begin{aligned} u = & \ln c - \frac{1}{\nu_f} \cdot \frac{h_f^{1+1/\psi}}{1+1/\psi} - \frac{1}{\nu_m} \cdot \frac{h_m^{1+1/\psi}}{1+1/\psi} \\ & + \frac{1}{\gamma_f} \cdot \varepsilon_f(k_f) + \frac{1}{\gamma_m} \cdot \varepsilon_m(k_m) + \frac{1}{\gamma_f} \cdot \epsilon_f(y_f, z_f) + \frac{1}{\gamma_m} \cdot \epsilon_f(y_m, z_m), \end{aligned} \tag{15}$$

which also includes the direct utility workers achieve from their horizontal and vertical labor-market choices, ϵ_g. In (15), ψ is the Frisch elasticity of labor supply to the market and ν_f and ν_m are the inverse weights on disutility from work.[24] Gender differences in ν, $\Delta\nu = \ln \nu_m - \ln \nu_f \neq 0$, measure differences in households willingness to supply the labor of men and women. This willingness can

[24] I consider as baseline a case of a homogeneous Frisch elasticity. I allow for gender differences in the Frisch elasticity in Appendix C.3

depend on the bargaining power of the spouses, their preferences, their productivity in non-market work, and also on social norms regarding gender roles.

Since the quantitative model has an intensive margin at which labor supply reacts to wage rates, firms take this into account when deciding upon wage offers. I assume that firms compete for workers in terms of short-run wages. Put differently, an equilibrium is a situation where a unilateral deviation in pay for one period does not pay off to the respective firm.[25] In particular, the decisive elasticity for their wage setting is now the elasticity of total hours supplied to the respective firm,

$$\phi = \eta + \psi.$$

The Frisch elasticity governs labor-supply reactions to short-run wage fluctuations because households live forever and have unlimited access to a risk-free bond with interest rate i. Accordingly, the budget constraint of a couple household now is $c + b' = w_f\left(k_f\right)h_f + w_m\left(k_m\right)h_m + (1+i)b$, where a prime (') denotes next period. The intertemporal preferences of a household are described by maximizing $U = u + \beta U'$. Apart from determining the labor-supply reaction to transitory (off-equilibrium) wage changes, the intertemporal dimension of the model does not matter for most of the subsequent analysis.

Singles. Including singles is important because my proposed mechanism is absent for singles and ignoring singles may hence lead to an overestimation of the mechanism's importance. In the quantitative model, there is a fraction s of agents who is single. The remaining fraction $1-s$ lives in couples and acts as in the basic model. Singles do not have a partner but are otherwise identical to spouses in couples. A single of gender g has period utility

$$u_g = \ln c_g - \frac{2}{\nu_g} \cdot \frac{h_g^{1+1/\psi}}{1+1/\psi} + \frac{2}{\gamma_g} \cdot \varepsilon_g(k_g) + \frac{2}{\gamma_g} \cdot \epsilon_g(y_g, z_g) \tag{16}$$

and acts subject to $w_g = w_g(k_g)$ and $c_g + b'_g = w_g \cdot h_g + (1+i)b_g$. The factors 2 in the preference weights reflect that singles assign full weight to their own labor disutility and non-pay job utility while couple households only weigh both factors by 50% for each spouse. Hence, couples maximize the average utility of their members, see (15). This does not affect consumption, as it is public in the household.

[25] In Appendix C.3, I also consider the case of competition in permanent wages which delivers similar results.

Additional assumptions. Apart from non-pay job preferences, I consider all heterogeneity to be binary such that there are $2^7 = 128$ combinations of a married worker's gender g, productivity x, horizontal position y, and vertical position z as well as productivity $-a$, horizontal position $-y$, and vertical position $-z$ of the partner. Singles simply differ in their own gender, productivity, horizontal position, and vertical position such that there $2^4 = 16$ different groups of singles. There is hence a total of $2^7 + 2^4 = 144$ different groups of workers on the labor market. Within these groups, workers differ in their preferences over workplaces as described in Section 2.

As a technical assumption, I assume that workers take their vertical and horizontal choices before learning specific firms' non-pay job characteristics within the segments. This rules out that workers choose a particular segment of the labor market because there is a particular firm in this segment with very likeable job characteristics. Instead, workers take into account the expected utility from entering a specific segment which they can assess from the information they have about the segment. Further, this assumption implies that individual firms cannot attract specific workers into their segment. Rather, it is overall pay in the segment which influences workers' choices of segment. In line with Section 2, I assume that firms do not internalize the effect of their individual pay on the size of their segment.

I assume that firms within a segment can observe a worker's gender and marital status as well as which segment the partner works in but not the exact preferences over workplaces (as in Section 2). This implies that firms can condition pay on marital status and on partner characteristics. Hence, there are 144 different wage rates in the economy, one for each of the 144 cells discussed above. The structure of the model implies that the equilibrium wage rate within each cell does not depend on cell size which allows me to solve for the within-cell equilibrium and the selection of individuals into the cells separately.

In order to maintain this important property, I leave out from the quantitative model a number of aspects that are further discussed in the context of imperfect competition on the labor market and gender differences in labor market outcomes. Instead, I address them in a model version that does not feature within-gender heterogeneity apart from non-pay job preferences. Appendix C.3 presents further model extensions that include gender differences in the elasticity of labor supply to the market, home production, firm entry, and alternative forms of wage competition between firms. I perform the main counterfactuals also within these model versions and compare results to those obtained from the baseline model with and without within-gender heterogeneity. The results show that, while some of the aspects affect the *level* of, e.g., wage rates to a non-negligible degree, gender *gaps*, which are the focus of this paper, behave very similarly across model versions which justifies abstracting from the discussed aspects in the main quantitative model.

Formal summary of the quantitative model. The quantitative model consists of 144 cells and households' self-selection into these cells. A cell is defined by worker's gender, marital status, and productivity, their choice of vertical and horizontal labor-market segment, and - for married individuals - productivity and vertical and horizontal labor-market segment of the partner.

The within-cell steady-state equilibrium is described as follows. In cell $p = \{g, x, y, z, -x, -y, -z\}$ which includes married individuals with gender g, productivity x, vertical position y, and horizontal position z, as well as partner's productivity $-x$, partner's vertical position $-y$, and partner's horizontal position $-z$, the following conditions describe the steady state,

$$w_p = a_x \cdot \alpha_y \cdot \frac{1}{1 + 1/\phi_p}, \tag{17}$$

$$\phi_p = \gamma_g \cdot V_z \cdot w_p h_p / (w_p h_p + w_{-p} h_{-p}) + \psi, \tag{18}$$

$$h_p^{1/\psi} = \nu_g \cdot c_p^{-1} \cdot w_p, \tag{19}$$

$$c_p = w_p h_p + w_{-p} h_{-p}, \tag{20}$$

where $-p = \{-g, -x, -y, -z, x, y, z\}$ describes the partner's cell. Conditions (17) to (20) describe, respectively, firms' profit-maximizing wage offers, the elasticity of total hours worked (combining extensive and intensive margin), the first-order condition for hours worked, and the steady-state household budget constraint.

In cell $n = \{g, x, y, z\}$ that includes single individuals with gender g, productivity x, vertical position y, and horizontal position z, the steady state is described by

$$
\begin{aligned}
w_n &= a_x \cdot \alpha_y \cdot \frac{1}{1+1/\phi_n}, \\
\phi_n &= \gamma_g/2 \cdot V_z + \psi, \\
h_n^{1\psi} &= \nu_g/2 \cdot c_n^{-1} \cdot w_n, \\
c_n &= w_n h_n.
\end{aligned}
$$

Type-I EV distributed taste shifters allow to determine cell choices as follows. Among a group of married individuals with exogenous characteristics g, x, and $-x$, feasible cell combination p', $-p'$ is chosen by share

$$
s_{p',-p'} = \frac{\exp(\widetilde{u}_{p',-p'}/\sigma)}{\sum_P \exp(\widetilde{u}_{p,-p}/\sigma)},
$$

where $\widetilde{u}_p$ is household utility (as described by (15)) net of taste shifters and P is the set of feasible choices for the considered type of couple. Analogously, among a group of single individuals with exogenous characteristics g and x, feasible cell p is chosen by share

$$
s_n = \frac{\exp(\widetilde{u}_n/\sigma)}{\sum_N \exp(u_n/\sigma)},
$$

where $\widetilde{U}_n$ is utility (as described by (16)) net of taste shifters and N is the set of feasible choices for the group.[26] The overall share of workers in a particular cell is obtained by multiplying s_n and s_p respectively, with the share of workers that have the particular characteristics which gives $s_p \cdot (1-s) \cdot s_g^x \cdot s_{-g}^{-x}$ and $s_n \cdot s \cdot s_g^x$ where s is the singles share and s_g^x is the share of workers with gender g who have productivity x.

4.2 Calibration

I now describe the parametrization of my model which is a combination of setting some parameters and calibrating others. The calibration targets the present-day U.S. economy and is described in Table 8. I set the share of couple households to 0.7 which is its 2015 counterpart in the data. The Frisch elasticity of labor supply to the market is set to 0.5 in accordance with Domeij and

[26]It is sufficient to consider the expected value of period utility here since also the direct utility gains or losses from choosing a specific segment accrue every period.

Flodén (2006).[27] The time preference rate is set to 0.98 to achieve a two percent real interest rate, interpreting a period as one year. The share of workers for whom preferences are redrawn each period is set to 0.2 mimicking a separation rate of 20%.[28] I consider an AR(1) process for firm-specific log productivity and take its persistence and standard deviation (0.97 and 0.09, respectively) from Bachmann and Bayer (2009).[29]

Gender-specific utility weights and productivity distributions are chosen to match the BLS estimates for the gender gap in average wage rates and average earnings for 2015, $\Delta w = 19.9$ log points and $\Delta wh = 32.6$ log points, and gender-specific average elasticities of labor supply to individual firms of $\eta_f = 1.793$ and $\eta_m = 2.413$ which are the estimates from Ransom and Oaxaca (2010) for their most recent sample. Normalizing the average female wage rate to $\bar{w}_f = 1$ and average hours worked to $\bar{h} = 0.33$, the targeted gaps imply $\bar{w}_m = 1.22$, $\bar{h}_f = 0.31$, and $\bar{h}_m = 0.35$. I normalize the low productivity level to 1 and the male share with high productivity to 0.5 and achieve the targets by setting the high productivity level to 2.06, the female high productivity share to 0.32, and the inverse labor disutility weights to $\nu_f = 0.038$ and $\nu_m = 0.041$. The average elasticities of labor supply to individual firms are matched by setting the inverse utility weights on non-pay job attributes to $\gamma_f = 0.417$ and $\gamma_m = 0.441$.

I proceed as follows to calibrate the parameters that govern vertical and horizontal differences between labor market segments interpreting the horizontal dimension of the labor market as industries with varying employer concentration and the vertical dimension as college education with differences in productivity. I first normalize the average value of the scale parameter of the taste shifters to 10 and the average productivity shifter to 1. I set the marginal productivity shifters to two thirds and four thirds which mimics a college to no-college wage ratio of 2 while pertaining an average value of one.[30] I set the scale parameters of the taste shifters to 5 and 15 implying that industry 2 is three times as concentrated as industry 1 in line with most recent observation (2012) of the 80-20 ratio of the C4 concentration index across three-digit manufacturing industries

[27] I eschew gender differences in the Frisch elasticities. This is due to a combination of two points. First, I want to limit the dimensions of exogenous gender differences for the counterfactuals. Second, empirical gender differences in Frisch elasticities are likely limited once estimation biases are accounted for (Bredemeier et al. 2019). Note that the model features endogenous gender differences in uncompensated (Marshallian) labor-supply elasticities due to weaker income effects for married women compared to married men.

[28] I use annual quit rates for the total economy from the BLS, see https://www.bls.gov/news.release/jolts.t18.htm.

[29] I vary the parameters of the process for firm-specific productivity in Appendix C.1.

[30] In 2015, the average hourly wage rate of workers with at most a high-school degree was \$16.96 while it was \$34.07 for workers with a bachelor degree or more (own aggregation based on Valletta, 2018, Table 2).

Table 8: Parameter values.

Symbol	Interpretation	Value	Target/Source
Aggregate parameters			
s	share of couples	0.700	observed
ψ	Frisch elasticity (to market)	0.500	Domeij and Flodén (2006)
σ	scale parameter taste shifters	0.516	high-school rate
β	time preference rate	0.980	real interest rate
θ	share re-drawn job preferences	0.111	average quit rate
ρ	persistence firm-productivity shocks	0.995	Bachmann and Bayer (2009)
σ_a	std. dev. firm-productivity shocks	0.12	Bachmann and Bayer (2009)
Parameters governing within-gender heterogeneity			
a_x	productivity level		
	high	2.055	normalize
	low	1.000	$\bar{w}_f = 1$
α_y	marginal revenue product shifter		
	high position	1.333	college
	low position	0.667	wage premium
V_z	(inverse) firm concentration		
	industry 1	1.50	80-20 ratio
	industry 2	0.50	firm concentration
Parameters governing exogenous gender differences			
s_g^{high}	share of workers with high productivity		
	female	0.320	wage
	male	0.500	gap
ν_g	inv. utility weight on labor supply		
	female	0.038	hours
	male	0.041	gap
γ_g	inv. utility weight, non-pay attributes		
	female	0.404	gap in inter-
	male	0.430	firm mobility
Implied gender gaps in exogenous factors			
$\Delta\bar{a}$	gap in average productivity	13.3 lp	$\Delta w = 19.9\,\text{lp}$
$\Delta\nu$	gap in labor disutility	9.9 lp	$\Delta h = 12.7\,\text{lp}$
$\Delta\gamma$	gap in weight on non-pay attributes	6.3 lp	$\Delta\eta = 29.7\,\text{lp}$

Notes: $\Delta x = \ln x_m - \ln x_m$. lp = log points.

while pertaining an average value of 1.[31] I set the variance of the taste shifters to 0.52 to match an equilibrium share of workers with the high vertical position of two thirds corresponding to the 2015 share of people with more than a high-school degree in the labor force.[32]

[31] The C4 index is the market share of the four largest firms in an industry. It is provided by the Census Bureau under the North American Industry Classification System (NAICS) for 1997, 2002, 2007, and 2012. While differences across industries in concentration are important for the model, its level simply affects the calibration of γ_g, see (23).

[32] This number stems from the BLS Spotlight on Statistics "Profile Of The Labor Force By Educational Attainment" from August 2017.

Table 9: Non-targeted moments.

	model	data
married to single wage gap		
among men	5 lp	6 lp
among women	-12 lp	-5 lp
gender wage gap		
among married	25 lp	25 lp
among singles	8 lp	9 lp
part-time to full-time gap		
earnings, among men	138 lp	118 lp
earnings, among women	135 lp	98 lp
wage rates, among men	88 lp	82 lp
wage rates, among women	86 lp	64 lp

Notes: Married-single gaps and gender gaps by marital status (age controlled) from Killewald and Lundberg (2017) and Cheng (2016). Part-time to full-time gaps by population group calculated from the BLS's 'Labor Statistics from the CPS', Tables 37 and 38 and the BLS's chart 'Time spent working by full- and part-time status, gender, and location'. In the model, I define part time in a way identical to the BLS definition, i.e., working 35 hours or less per week where I define $\bar{h}_m$ to correspond to 40 hours per week. lp = log points.

4.3 Status-quo analysis

It is worthwhile to compare gender gaps in exogenous factors to gender gaps in endogenous labor-market outcomes. The model does not need large exogenous differences between genders, which are both state at the bottom of Table 8. Specifically, the gender gap in average productivity is only about 13 log points while the gender gap in wage rates, which amplifies the productivity gap through monopsonistic discrimination and gender segregation, is about 20 log points. The gender gap in the preference weight on labor supply which influences the gender gap in hours worked is only about 10% while men work 14% longer hours than do women. Most importantly, I need to put into the model only very little deep gender differences in the importance of non-pay job characteristics. The gender gap in the exogenous utility parameter γ is only about 6% while men are roughly a third more mobile between firms than women. Put differently, only about one fifth of the gender gap in inter-firm mobility the model devotes to exogenous factors while over 80% of the gap are explained endogenously within the model.

Table 9 considers non-targeted moments. While the gender wage gap is a targeted moment, its distribution is non-targeted. In line with the data, my model generates marriage wage premia for

Table 10: The wage sensitivity of quits in artificial data from simulated model.

	(1) men	(2) women	(3) men	(4) women
log wage rate	-0.0488^{***}	-0.0285^{***}	-0.0661^{***}	-0.0658^{***}
	(0.0034)	(0.0031)	(0.0037)	(0.0043)
log wage rate × (earnings share-0.5)			-0.1404^{***}	-0.1481^{***}
			(0.0106)	(0.0115)
observations	20197	20197	20197	20197

Notes: Mean estimates and standard errors from 10,000 Monte-Carlo repetitions. Dependent variable: quit between observation year and following year; linear probability models; standard errors in parentheses; *, **, *** indicate $p < 0.1$, $p < 0.05$, $p < 0.01$; alternative wage determinant (X): mean wage in respective cell p; additionally included variables: constant, earnings share.

men and marriage wage penalties for women. The reason is that particularly couples are subject to the proposed mechanism that leads to an endogenous gender gap in inter-firm mobility. In line with the findings by Webber (2016), the gender gap in inter-firm mobility is sufficiently more pronounced among married individuals than among singles. Quantitatively, the model generates marriage premia for men and marriage penalties for women of a magnitude that is quite reasonable compared to the data and, accordingly, the gender wage gap is smaller among singles and larger among married individuals to a data-consistent degree, see Table 9. The model also features part-time penalties as the wish to work shorter hours coincides with a particular earner role within the family that makes the specific worker rather immobile between firms. Firms can exploit this by paying less to workers who tend to work shorter hours. Quantitatively, empirical part-time penalties are matched quite successfully by the model, see the lower part of Table 9. The good model performance with respect to non-targeted moments provides confidence that the calibrated model is a suitable laboratory for counterfactual analysis.

Before turning to counterfactual analysis within the model, I use the model for a final status-quo analysis and perform quit regressions like the one considered in the empirical analysis. Specifically, I simulate the model around its steady state, generate artificial data which I then use to perform the same estimations as in Table 4 for the real-world data. The aim of this exercise is twofold. First, I want to make sure that the model features relations between quits and wages that are in line with their empirical counterparts because this relation is the moment that identifies gender differences in the elasticity of labor supply in many empirical studies. Second, I want to corroborate the ability of my extended quit regressions to differentiate differences in household earner roles from deep gender differences in the importance of non-pay job characteristics.

I describe the simulation and estimation in detail in Appendix C.1 and concentrate on the results here. Table 10 shows mean results from 10,000 Monte-Carlo estimations from the simulated data. Qualitatively, the results resemble the empirical ones presented in Table 4. Also for the simulated model data, standard quit regressions without taking into account household earner roles suggest that men are considerably more mobile between firms than women. Hence, quit regressions are able to detect gender differences in inter-firm mobility in the simulated data. Further, the size of the coefficients in the first two columns of Table 10 are of similar magnitude to the ones in Table 4 indicating that the model features a realistic relation between quits and wage rates.

When I include the interaction between the contributed earnings share and the log wage rate, I obtain negative coefficients in the Monte-Carlo lab as I did in the simplified environment of the basic model and in the real-world data. In Appendix C.1, I corroborate this finding for alternative calibrations of the firm-specific productivity process. While estimates of the coefficients on the interaction terms are somewhat larger than in the empirical results in Table 4, the remaining gender difference in the coefficient on the non-interacted log wage rates are small as are the exogenously fed-in gender differences in the importance of non-pay job characteristics.

4.4 Counterfactual analysis

I perform two series of counterfactuals which aim at determining the importance of the different gender gaps in exogenous factors to the endogenous gender gaps in labor-market outcomes and at quantifying the degree of amplification due to the feedback mechanism between the gender gaps.

Decomposition. In order to decompose the gender gap in labor-market outcomes into the consequence of the three exogenous gender differences that the model knows, gender differences in labor-demand factors, in labor-supply (preference) factors, and in the importance of non-pay job characteristics (or deep mobility differences), I perform counterfactual model evaluations where I shut of gender differences in preferences by setting the respective parameter for women to its baseline value for men. Specifically, I first shut off gender differences in the disutility of work, i.e., set $\Delta\nu$ to zero. Second, I shut off gender differences in the utility weights on non-pay job attributes,

Table 11: Shutting off gender differences in preferences.

	full model (endogenous mobility)			**reference model** (exogenous mobility)		
	Δw	Δwh	$\Delta \eta$	Δw	Δwh	$\Delta \eta$
A) Full calibration	19.9 lp	32.6 lp	29.7 lp	19.9 lp	32.6 lp	29.7 lp
B) $\Delta\nu = 0$	18.6 lp	26.3 lp	25.7 lp	19.5 lp	26.9 lp	29.7 lp
C) $\Delta\gamma = 0$	17.8 lp	29.9 lp	22.1 lp	12.6 lp	22.9 lp	0.0 lp
D) $\Delta\nu = 0$ and $\Delta\gamma = 0$	16.6 lp	23.6 lp	17.2 lp	12.3 lp	17.2 lp	0.0 lp
Relative contribution of ...						
... $\Delta\bar{a}$ $(= D/A)$	83.8%	72.4%	57.9%	62.1%	52.8%	0.0%
... $\Delta\gamma$ $(= (B - D)/A)$	10.1%	8.3%	28.6%	36.4%	29.8%	100.0%
... $\Delta\nu$ $(= (C - D)/A)$	6.1%	19.3%	13.1%	1.5%	17.5%	0.0%

Notes: $\Delta x = \ln x_m - \ln x_m$. $\Delta x = 0$ means x_f is set to x_m. lp = log points. Levels of wage rates, hours, and elasticities of labor supply to individual firms shown in Table 15 in Appendix C.2.

i.e., set $\Delta\gamma$ to zero. Third, I shut off all gender differences in preferences, i.e., set both $\Delta\nu$ and $\Delta\gamma$ to zero. I do this for my full model and for a reference model in which I shut off my mechanism and treat inter-firm mobility as exogenous, i.e., a model version where η is a parameter.[33]

The results are shown in the upper block of Table 11 where the left block refers to my full model and the right block refers to the reference model with exogenous mobility. The numbers in the table give the gender gaps in wage rates, earnings, and inter-firm mobility in the different counterfactuals (in log points). The lower block of the Table 11 shows the percentage contributions of the different exogenous gender gaps on the endogenous gaps, which are obtained based on the counterfactual evaluations.

The results show that my full model with endogenous inter-firm mobility generates substantial gender gaps in wage rates, earnings, and inter-firm mobility also without gender differences in preferences. Without any gender differences in preferences, i.e., with gender differences only in labor-demand factors, captured in the model through differences in marginal revenue products, the model still generates almost 85% of the gender wage gap, almost 75% of the gender gap in earnings, and more than 60% of the gender gap in inter-firm mobility. The reference model with exogenous inter-firm mobility assigns substantially less importance to gender differences in labor-demand factors.

Reversely, the contribution of gender gaps in preferences to gender gaps in labor-market out-

[33] To focus on inter-firm mobility, I continue to model the choices of labor-market segments as a joint decision of couples.

comes is limited in my full model. In particular, the gender gap in the importance of non-pay job characteristics contributes only 10% to the gender wage gap and, strikingly, only less than 30% to the gender gap in inter-firm mobility.[34] In the reference model with exogenous inter-firm mobility, by contrast, gender differences in the importance of non-pay job characteristics explain more than 35% of the gender wage gap and, by construction, 100% of the gender gap in inter-firm mobility. Put differently, a model without endogenous inter-firm mobility hugely overestimates the causal role of deep gender differences in preferences over jobs.

As discussed in Section 2, my model suggests that gender differences in labor-supply factors have an influence on how mobile men and women are between firms and how strongly firms can discriminate monopsonistically against women. The quantitative results in Table 11 show that this effect is not negligible as about 6% of the gender wage gap is assigned to gender differences in labor-supply factors. This effect is overlooked by a model where inter-firm mobility is considered as exogenous. In the reference model, labor-supply factors only exert a small effect on gender-specific wages through composition effects that reflect the segregation across labor-market segments.

Amplification. In my second set of counterfactuals, I vary parameters in order to assess the quantitative degree of amplification due to endogenous inter-firm mobility. Specifically, I close one of the gender gaps in exogenous parameters in each experiment. Put differently, I make men and women more similar in the respective dimension. The results are shown in Table 12. The numbers give the gender gaps in wage rates, earnings, and inter-firm mobility for the baseline calibration and the different counterfactuals (in log points). Numbers in brackets give the absolute change (in log points) in the respective gap relative to the baseline calibration with all three exogenous gender gaps shut on. As in Table 11, the left block shows the results for my full model with endogenous inter-firm mobility while the right block shows the results for the reference model where inter-firm mobility is exogenous.

First, I raise firms' demand for female labor by raising women's productivity by ten log points. In the reference model with exogenous inter-firm mobility, this closes the gender wage gap by little more than ten log points and, by construction, leaves the gender gap in inter-firm mobility

[34] The latter number is larger than the 20% which results from a direct comparison of $\Delta\gamma$ and $\Delta\eta$ because the (small) $\Delta\gamma$ generates some endogenous earnings gap and thus also exerts an indirect, general-equilibrium effect on $\Delta\eta$ in my full model.

Table 12: Closing exogenous gender gaps.

	full model (endogenous mobility)			reference model (exogenous mobility)		
	Δw	Δwh	$\Delta \eta$	Δw	Δwh	$\Delta \eta$
baseline calibration	19.9 lp	32.6 lp	29.7 lp	19.9 lp	32.6 lp	29.7 lp
change in demand	6.0 lp	13.5 lp	17.6 lp	9.0 lp	18.1 lp	29.7 lp
$\Delta a \downarrow 10\,\text{lp}$	(-13.9 lp)	(-19.1 lp)	(-12.1 lp)	(-10.8 lp)	(-14.5 lp)	(±0.0 lp)
change in supply	18.6 lp	26.2 lp	25.7 lp	19.6 lp	27.8 lp	29.7 lp
$\Delta \nu \downarrow 10\,\text{lp}$	(-1.3 lp)	(-6.4 lp)	(-4.0 lp)	(-0.2 lp)	(-4.8 lp)	(±0.0)
change in mobility	16.6 lp	28.2 lp	15.8 lp	17.3 lp	29.2 lp	19.7 lp
$\Delta \gamma \downarrow 10\,\text{lp}$	(-3.2 lp)	(-4.4 lp)	(-13.9 lp)	(-2.5 lp)	(-3.4 lp)	(-10.0 lp)

Notes: Closures of exogenous gaps achieved through changing female parameter accordingly. $\Delta x = \ln x_m - \ln x_f$. lp = log points. Levels of wage rates, hours, and elasticities of labor supply to individual firms shown in Table 15 in Appendix C.2.

untouched. By contrast, in my full model, the effects are amplified through the mutually enforcing cycle between gender gaps in pay and inter-firm mobility. The cycle amplifies the effect on the gender wage gap by about 30% as the gap closes by about 14 log points in my full model. The reduction in the gender gap in inter-firm mobility is also considerable which reduces firms' ability to discriminate monopsonistically against women.

Second, I reduce the gender gap in the exogenous utility weight on labor supply inducing households to increase female labor supply. In the model with exogenous mobility, this has only a very small effect on the gender wage gap stemming from more women choosing better paying labor-market segments. Again, the gender gap in inter-firm mobility is unaffected by construction. Also here, effects are substantially stronger in my full model with endogenous inter-firm mobility. The reduction in the gender wage gap is more than six times as strong due to the endogenous increase in women's relative inter-firm mobility which raises their pay.

Third, I reduce women's exogenous preference weight on non-pay job attributes which makes them more mobile between firms. In my model, this exogenous impulse is amplified through the mutually reinforcing cycle and the gender gap in inter-firm mobility is reduced by about 14 log points. This also leads to a quite substantial reduction in the gender wage gap by more than 3 log points. Effects are substantially smaller in the reference model.

The counterfactual changes in the exogenous gender gaps can be used to think about the effects of policy in my model. A change in $\Delta\bar{a}$ can be understood as a stand-in for anything that induces a movement along the relative labor-supply curve of men versus women. Examples could be payroll subsidies that differ by gender of the worker or an elimination of the high marginal tax rates imposed on secondary earners under progressive joint taxation. Similarly, a change in $\Delta\nu$ can be thought of as anything that pushes the relative labor-supply curves, hence lets households change the proportion at which they supply male and female labor for given relative wage rates. Policy examples that may raise female labor supply could be, for example, extended access to child care and reforms of divorce or alimony legislation. Finally, there can be policies that effectively change $\Delta\gamma$ through making firms more similar in certain dimensions of non-pay job attributes. For example, if allowing child-sick leave is mandatory for firms, this is one less dimension in which firms differ and this may make especially women more mobile between firms. In my model, the effects of such policies are amplified by reinforcing changes in relative wage mark-downs imposed on men and women by monopsonistic firms. Put differently, my analysis suggests that such policies have substantially stronger effects on gender gaps in labor-market outcomes than one would expect if one neglected the endogeneity of inter-firm mobility.

5 Conclusion

In this paper, I argue that the gender gap in inter-firm mobility is largely due to men's and women's different earner roles within households, rather than intrinsic differences between genders. This relation stems from a structural model where households decide endogenously how important pay and non-pay characteristics are for the job choices of their members. I have presented direct empirical evidence on the role of household earner roles for the quitting behavior of workers that supports my theoretical prediction. Quantitative model evaluations suggest that the endogeneity of inter-firm mobility is important. If one mistakes inter-firm mobility as exogenous, one underestimates the role of labor-demand and labor-supply factors for the gender wage gap, overestimates the role of gender differences in the importance of different job attributes, and underestimates the effects of gender-targeted policy reforms and changes in social norms.

It is interesting to consider two further policy implications of my results. First, it is well established that both minimum wages and payroll subsidies can be appropriate policy tools to achieve efficiency on monopsonistic labor markets, i.e., to equate the marginal product of labor with the marginal rate of substitution between consumption and leisure. Heterogeneity in marginal products across workers is an argument against a uniform minimum wage. My analysis suggests that monopsonistic mark-downs – and hence the required subsidy rates – are likely heterogeneous as well and depend on, among other things, a worker's family background. This implies a trade-off between heterogenous marginal products and heterogenous wage mark-downs and it is a priori not clear whether a uniform minimum wage or a uniform subsidy rate induces the larger efficiency loss.

Second, one-wage or "equal pay for equal work" policies are frequently discussed to fight monopsonistic discrimination against women. In presence of gender segregation across labor-market segments formed by, e.g., industries and occupations, however, the equalizing effects of such policies on gender-specific pay are limited. Firms would then base wage offers on the average worker mobility in their specific segment of the labor market such that the effect on female workers in, mostly low-pay, majority-female segments would be small but the effect on female workers in, mostly higher-pay, majority-male segments would be large. This suggests a policy trade-off between raising female pay on average and increasing income inequality between women. Both these policy trade-offs constitute interesting avenues for future research.

References

Aizer, A. (2010). The gender wage gap and domestic violence. *American Economic Review 100*(4), 1847–1859.

Alesina, A., A. Ichino, and L. Karabarbounis (2011). Gender-based taxation and the division of family chores. *American Economic Journal: Economic Policy 3*, 1–40.

Artz, B., A. H. Goodall, and A. J. Oswald (2018). Do women ask? *Industrial Relations 57*(4), 611–636.

Autor, D., A. Kostøol, M. Mogstad, and B. Setzler (2019). Disability benefits, consumption insurance, and household labor supply. *American Economic Review 109*(7), 2613–54.

Babcock, L. and S. Laschever (2009). *Women don't ask: Negotiation and the gender divide.* Princeton University Press.

Babcock, L., M. P. Recalde, L. Vesterlund, and L. Weingart (2017). Gender differences in accepting and receiving requests for tasks with low promotability. *American Economic Review 107*(3), 714–747.

Bachmann, R. and C. Bayer (2009). Firm-specific productivity risk over the business cycle: Facts and aggregate implications. CESifo Working Paper Series 2844.

Bachmann, R. and C. Bayer (2014). Investment dispersion and the business cycle. *American Economic Review 104*(4), 1392–1416.

Bachmann, R., R. J. Caballero, and E. M. R. A. Engel (2013). Aggregate implications of lumpy investment: New evidence and a DSGE model. *American Economic Journal: Macroeconomics 5*(4), 29–67.

Bailey, M. J., B. Hershbein, and A. R. Miller (2012). The opt-in revolution? Contraception and the gender gap in wages. *American Economic Journal: Applied Economics 4*(3), 225–254.

Barth, E. and H. Dale-Olsen (2009). Monopsonistic discrimination, worker turnover, and the gender wage gap. *Labour Economics 16*(5), 589–597.

Bayard, K., J. Hellerstein, D. Neumark, and K. Troske (2003). New evidence on sex segregation and sex differences in wages from matched employee-employer data. *Journal of Labor Economics 21*(4), 887–922.

Beaudry, P. and E. Lewis (2014). Do male-female wage differentials reflect differences in the return to skill? Cross-city evidence from 1980-2000. *American Economic Journal: Applied Economics 6*(2), 178–194.

Bhashkar, V. and T. To (2003). Oligopsony and the distribution of wages. *European Economic Review 47*(2), 371–399.

Bick, A. (2016). The quantitative role of child care for female labor force participation and fertility. *Journal of the European Economic Association 14*(3), 639–668.

Blackaby, D., A. L. Booth, and J. Frank (2005). Outside offers and the gender pay gap: Empirical evidence from the UK academic labour market. *Economic Journal 115*(501), 81–107.

Blau, F. D. and L. M. Kahn (2017). The gender wage gap: Extent, trends, and explanations.

Journal of Economic Literature 55(3), 789–865.

Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2016). Consumption inequality and family labor supply. *American Economic Review 106*(2), 387–435.

Booth, A. L., M. Francesconi, and J. Frank (2003). A sticky floors model of promotion, pay, and gender. *European Economic Review 47*(2), 295–322.

Booth, A. L. and P. Katic (2011). Estimating the wage elasticity of labour supply to a firm: What evidence is there for monopsony? *Economic Record 87*(278), 359–369.

Bosquet, C., P.-P. Combes, and C. García-Peñalosa (2019). Gender and promotions: Evidence from academic economists in France. *The Scandinavian Journal of Economics forthcoming.*

Bredemeier, C., J. Gravert, and F. Jüßen (2019). Estimation labor-supply elasticities with joint borrowing constraints of couples. *Journal of Labor Economics forthcoming.*

Bredemeier, C. and F. Jüßen (2013). Assortative mating and female labor supply. *Journal of Labor Economics 31*(3), 603 – 631.

Card, D., A. R. Cardoso, J. Heining, and P. Kline (2018). Firms and labor market inequality: Evidence and some theory. *Journal of Labor Economics 36*(S1), 13–70.

Cheng, S. (2016). The accumulation of (dis)advantage: The intersection of gender and race in the long-term wage effect of marriage. *American Sociological Review 81*(1), 29–56.

Cherchye, L., B. De Rock, and F. Vermeulen (2012). Married with children: A collective labor supply model with detailed time use and intrahousehold expenditure information. *American Economic Review 102*(7), 3377–3405.

Croson, R. and U. Gneezy (2009). Gender differences in preferences. *Journal of Economic Literature 47*(2), 448–474.

Depew, B. and T. A. Sørensen (2013). The elasticity of labor supply to the firm over the business cycle. *Labour Economics 24*(C), 196–204.

Domeij, D. and M. Flodén (2006). The labor-supply elasticity and borrowing constraints: Why estimates are biased. *Review of Economic Dynamics 9*, 242–262.

Donni, O. and N. Moreau (2007). Collective Labor Supply: A Single-Equation Model and Some Evidence from French Data. *Journal of Human Resources 42*(1), 214–246.

Félix, S. and P. Portugal (2017). Labor market imperfections and the firm's wage setting policy. Working Papers w201704, Banco de Portugal, Economics and Research Department.

Gobillon, L., D. Meurs, and S. Roux (2015). Estimating gender differences in access to jobs. *Journal of Labor Economics 33*(2), 317–363.

Goldin, C. (2014). A grand gender convergence: Its last chapter. *American Economic Review 104*(4), 1091–1119.

Goldin, C., S. P. Kerr, C. Olivetti, and E. Barth (2017). The expanding gender earnings gap: Evidence from the LEHD-2000 census. *American Economic Review 107*(5), 110–114.

Grund, C. (2015). Gender pay gaps among highly educated professionals - compensation components do matter. *Labour Economics 34*(C), 118–126.

Heinz, M., H.-T. Normann, and H. A. Rau (2016). How competitiveness may cause a gender wage gap: Experimental evidence. *European Economic Review 90*, 336 – 349.

Hirsch, B. (2016). Gender wage discrimination - does the extent of competition in labor markets explain why female workers are paid less than men? *IZA World of Labor 310*, 1–10.

Hirsch, B. and E. J. Jahn (2015). Is there monopsonistic discrimination against immigrants? *Industrial and Labor Relations Review 68*(3), 501–528.

Hirsch, B., E. J. Jahn, A. Manning, and M. Oberfichtner (2019). The urban wage premium in imperfect labour markets. CEP Discussion Paper 1608.

Hirsch, B., E. J. Jahn, and C. Schnabel (2018). Do employers have more monopsony power in slack labor markets? *Industrial and Labor Relations Review 71*(3), 676–704.

Hirsch, B., T. Schank, and C. Schnabel (2010). 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 28*(2), 291–330.

Juhn, C. and K. McCue (2017). Specialization then and now: Marriage, children, and the gender earnings gap across cohorts. *Journal of Economic Perspectives 31*(1), 183–204.

Keane, M. P. (2011). Labor supply and taxes: A survey. *Journal of Economic Literature 49*(4), 961–1075.

Killewald, A. and I. Lundberg (2017). New evidence against a causal marriage wage premium. *Demography 54*(3), 1007–1028.

Manning, A. (2003). *Monopsony in Motion*. Princeton: Princeton University Press.

Manning, A. (2011). *Imperfect Competition in the Labor Market*, Volume 4 of *Handbook of Labor Economics*, Chapter 11, pp. 973–1041. Elsevier.

Manning, A. and F. Saidi (2010). Understanding the gender pay gap: What's competition got to do with it? *Industrial and Labor Relations Review 63*(4), 681–698.

Mas, A. and A. Pallais (2017). Valuing alternative work arrangements. *American Economic Review 107*(12).

McFadden, D. L. (1973). Conditional logit analysis of qualitative choice behavior. In P. Zarembka (Ed.), *Frontiers in econometrics*, pp. 105–142. New York: Academic Press.

McGee, A., P. McGee, and J. Pan (2015). Performance pay, competitiveness, and the gender wage gap: Evidence from the United States. *Economics Letters 128*(C), 35–38.

Nishiyama, S. (2019). The joint labor supply decision of married couples and the U.S. social security pension system. *Review of Economic Dynamics 31*, 277–304.

Ortigueira, S. and N. Siassi (2013). How important is intra-household risk sharing for savings and labor supply? *Journal of Monetary Economics 60*, 650–666.

Pekkarinen, T. and J. Vartiainen (2006). Gender differences in promotion on a job ladder: Evidence from Finnish metalworkers. *Industrial and Labor Relations Review 59*(2), 285–301.

Ransom, M. and R. L. Oaxaca (2005). Intrafirm mobility and sex differences in pay. *Industrial and Labor Relations Review 58*(2), 219–237.

Ransom, M. R. and R. L. Oaxaca (2010). New market power models and sex differences in pay. *Journal of Labor Economics 28*(2), 267–289.

Redmond, P. and S. McGuinness (2019). The gender wage gap in Europe: Job preferences, gender convergence and distributional effects. *Oxford Bulletin of Economics and Statistics 81*(3), 564–587.

Sulis, G. (2011). What can monopsony explain of the gender wage differential in Italy? *International Journal of Manpower 32*(4), 446–470.

Sullivan, P. and T. To (2014). Search and non-wage job characteristics. *Journal of Human Resources 49*(2), 472–507.

Valletta, R. G. (2018). Recent flattening in the higher education wage premium: Polarization, skill downgrading, or both? In *Education, Skills, and Technical Change: Implications for Future U.S. GDP Growth*, pp. 313–342. National Bureau of Economic Research.

Webber, D. A. (2016). Firm-level monopsony and the gender pay gap. *Industrial Relations 55*(2), 323–345.

Wiswall, M. and B. Zafar (2017). Preference for the workplace, investment in human capital, and gender. *The Quarterly Journal of Economics 133*(1), 457–507.

Appendix

A Additional analytical results

A.1 Dynamic monopsony representation

Here I, present the firm problem in the style of the dynamic monopsony model of the labor market (Manning, 2003). Inflow and outflow rates of workers of group g, i.e. quit rates and the number of recruits, depend on wage rates offered by the firm to workers of this group and are denoted by $q_g\left(w_g\right)$ and $r_g\left(w_g\right)$, respectively, where I suppress a firm index for convenience. It is an important contribution of my paper that I have derived the functions $q_g(w_g)$ and $r_g(w_g)$ as endogenous equilibrium objects in Section 2.6.

Group-specific employment at a firm evolves according to

$$n_{g,t} = \left(1 - q_g\left(w_{g,t}\right)\right) n_{g,t-1} + r_g\left(w_{g,t}\right).$$

In a steady state where the number of quits equals the number of recruits, $q\left(w_g\right) = r_g\left(w_g\right)$, labor supplied by group g to the considered firm is given by

$$n_g = \frac{r_g\left(w_g\right)}{q_g\left(w_g\right)}.$$

Steady-state profits obtained from employing workers of group g are

$$\left(a_g - w_g\right) \cdot \frac{r_g\left(w_g\right)}{q_g\left(w_g\right)} \cdot h_g,$$

where a_g is the marginal revenue product and h_g are hours worked per worker. Profits are maximized by wage offer

$$w_g = \alpha_g \cdot \frac{1}{1 + 1/\eta_g} \tag{21}$$

where

$$\eta_g = \frac{r_g' w_g}{r_g} - \frac{q_g' w_g}{q_g}$$

is the elasticity of the number of workers who supply their labor to a given individual firms which is the sum of the (absolute) elasticities of recruits and the quit rate.

A.2 Alternative derivation of the elasticity of labor supply to individual firms

In this appendix, I derive the elasticity of labor supply to individual firms in a way that puts worker flows between firms and, in particular, workers' quit decisions into the center of the analysis. For this, I slightly change the way I model worker preferences over non-pay job characteristics and use a way similar to Bhashkar and To (2003). This change allows an easy representation of workers' firm choices in the spirit of the empirical analysis while it yields the same results regarding the elasticity of labor supply to individual firms as does the baseline model presented in Section 2.

Here, the non-pay characteristics of a job are mapped into a scalar v on the unit circle, like in a Salop model of product market competition. There are V firms with actual characteristics described by $v = 0,\ {}^{1}\!/\!{}_{V},\ {}^{2}\!/\!{}_{V},\ ...,\ 1$. Each worker is assigned a number which summarizes his or her ideal employer and these worker ideals are distributed uniformly on (0,1). Worker utility depends on the difference between his or her ideal workplace and the actual characteristics of the chosen employer. In particular, the household target function is

$$u = \ln c + \frac{1}{\gamma_f} \cdot (1 - |k_f - v_f|) + \frac{1}{\gamma_m} \cdot (1 - |k_m - v_m|)\,, \tag{22}$$

where k_g describes non-pay characteristics of the chosen firm, and v_g the worker's most preferred job characteristics.

To derive the elasticity of labor supply to individual firms, I consider the workers who work for a given firm and evaluate how many of these workers would decide not to work for this firm if the firm were to change its wage offer by factor z. Each individual worker bases his or her decision whether to continue to work for this firm or whether to switch firms on a comparison of the costs of staying with the firm and the costs Γ_g of switching to another firm,

$$z \cdot w_g \cdot h_g \cdot \lambda \gtrless \Gamma_g.$$

The left-hand side of the condition above describes the costs of staying with the firm and the right-hand side the costs of switching employers. The staying costs on the left-hand side reflect the absolute reduction in earnings, i.e., share z of labor earnings $w_g h_g$, which are translated into utility terms through multiplication with the marginal utility of wealth, i.e., the Lagrange parameter on the budget constraint, λ. The switching costs on the right-hand side capture the differences in utility from non-pay job characteristics between the current employer and the firm one would switch to (the utility value of a change in job attributes) and are here summarized by Γ_g which will be related to deep parameters of the model later. The share of workers who would switch to another firm is the share of workers for whom these quitting costs are less than $z \cdot w_g \cdot h_g \cdot \lambda$. The resulting number of workers who would remain at the firm after the pay cut is

$$n_g = \bar{n}_g \cdot \left(1 - \mathrm{cdf}_{\Gamma,g}\left(z \cdot w_g \cdot h_g \cdot \lambda\right)\right),$$

where $\bar{n}_g$ is the comparison employment level without the pay cut and $\mathrm{cdf}_{\Gamma,g}$ is the group-specific cumulated distribution function of quitting costs. Since z is the *relative* pay *cut*, the elasticity of labor supply to the individual firm is given by

$$\eta_g = -\frac{n_g}{z} \cdot \frac{1}{\bar{n}_g} = w_g \cdot h_g \cdot \lambda \cdot \mathrm{pdf}_{\Gamma,g}\left(z \cdot w_g \cdot h_g \cdot \lambda\right),$$

where $\mathrm{pdf}_{\Gamma,g}$ is the probability density function of switching costs among workers in the firm.

In a symmetric equilibrium, it is straightforward to determine the distribution of quitting costs Γ_g. The indifferent worker is located exactly in the middle between the two relevant firms and hence has zero non-pecuniary costs of switching between the two firms. On the other extreme, for the worker whose ideal workplace is exactly matched by the considered firm, switching to the next best firm is associated with a reduction of utility from non-pay job attributes by $1/(\gamma_g V)$. Hence, switching costs are distributed uniformly on $(0, {}^1\!/_{\gamma_g V})$ such that the density is constant and given by $\gamma_g V$. It follows that

$$\eta_g = w_g \cdot h_g \cdot \lambda \cdot \gamma_g \cdot V.$$

Combining this result with the budget constraint and the fact that marginal utility λ equals c^{-1} with log utility gives the elasticity of labor supply to individual firms by workers of gender g as

$$\eta_g = \gamma_g \cdot V \cdot \frac{w_g h_g}{w_g h_g + w_{-g} h_{-g}}. \tag{23}$$

Hence, this model version delivers the same result for η_g as the model presented in the main text.

A.3 Search costs

In my baseline model, there are no monetary quitting costs. Workers can react to pay cuts by instantaneously switching to another firm and potential losses in utility from non-pay job characteristics are the only cost of doing so. In reality, there are of course also monetary costs of quitting an employer, importantly search costs. After reentering employment, there are also earnings penalties for having been unemployed which increase in the length of the past unemployment spell. It is potentially important to consider these dimensions as women on average remain longer in unemployment than men. In this Appendix, I introduce an additional, fixed cost of quitting and show that this extension, while complicating the derivations, does not impact on gender-specific elasticities of labor supply to individual firms beyond a uniform rescaling of the preference parameters γ_f and γ_m. For convenience, I perform this extension within the framework set up in Appendix A.2 because, also here, I will focus on the quit decision of a worker.

I assume that, when quitting at an employer, an individual has to pay an additional (search) cost of Ω (expressed in utils) before being able to sign up at a new firm. I still allow workers to move to their (now) most preferred firm, so a way of interpreting Ω is as the cost of (perfectly directed) search. In the following, I suppress indices for convenience but Ω can be thought of being gender-specific. An individual now quits at a wage-cutting firm if

$$z \cdot w \cdot h \cdot \lambda > \Gamma + \Omega.$$

Hence, the quit rate is now $q = \mathrm{cdf}_{\Gamma+\Omega}\left(z \cdot wh \cdot \lambda\right)$ and the elasticity of labor supply to individual firms is

$$\eta = \frac{\partial q}{\partial z} = \mathrm{pdf}_{\Gamma+\Omega} \cdot wh \cdot \lambda.$$

Two things are worth noting about the above result. First, while Γ varies across individuals of a given gender, Ω is a fixed cost of quitting and hence a (gender-specific) constant. Second, in this version, the smallest value of Γ is negative in equilibrium. There are workers who have experienced a small change in their job preferences which would make them better off at a different firm but they stay at their previous employer to avoid the monetary quitting costs Ω. The indifferent worker prior to any pay cut is characterized through $\Gamma + \Omega = 0$ (rather than $\Gamma = 0$ as in the baseline model) such that the slightest pay cut z induces some workers to quit.

To determine the density function of total quitting costs, $\mathrm{pdf}_{\Gamma+\Omega}$, in equilibrium (where further pay cuts do not pay off for firms, i.e., at $z = 0$), one can concentrate on the group of workers for whom a marginal pay cut may induce quitting. These are the workers who would work for a different firm were it not for the fixed quitting costs Ω. For a firm j at location v in the job-characteristics space, these workers are located between $v - 1/(2V) - \gamma\Omega/2$ and $v - 1/(2V)$ as well as between $v + 1/(2V)$ and $v + 1/(2V) + \gamma\Omega/2$. In these intervals, workers would not work for the considered firm if Ω were zero but, for positive Ω, those who previously worked for the firm remain also after their job preferences have shifted into these intervals (share $1/V$ of the workers in these intervals). Thus, mass $\gamma\Omega/V$ of workers work for firm j because of the fixed quitting costs (share $\gamma\Omega$ of the firm's total workforce which is still $1/V$ in equilibrium). Within this group, the smallest value of total quitting costs is zero (the worker for whom $\Gamma = -\Omega$) and the largest is Ω (the worker for whom $\Gamma = 0$) and total quitting costs are distributed uniformly with density $1/\Omega$. Combined with share $\gamma\Omega$ of all workers at the firm falling in this group. The density of total quitting costs among all workers in the firm evaluated at zero is

$$\mathrm{pdf}_{\Gamma+\Omega}(0) = \gamma.$$

Hence, the equilibrium elasticity of labor supply to individual firms in this model version is

$$\eta = \gamma wh\lambda = \gamma whc^{-1}$$

which differs from its counterpart in the baseline model only in the absence of the constant V. It follows that in a calibrated version of this model version targeting gender-specific values for η, the values of the preferences parameter γ_f and γ_m would be rescaled but the gender gap in them, $\Delta\gamma = \ln\gamma_m - \ln\gamma_f$ would remain the same. Also any multiplicative change in γ_f as performed in the counterfactuals presented in Section 4 would have identical effects across model versions.

B Appendix to empirical analysis

B.1 Variable definitions

Quits. I first construct separations. A separation between years t and $t + 1$ occurs when either the individual was employed in year t but not in year $t + 1$ or the individual was employed in both years t and $t + 1$ and reports a tenure of one year or less in year $t + 1$. A quit is a voluntary separation which I define based on the answer to the question ”Why did your last job end?”. If the answer is ”quit; resigned; retired; pregnant; needed more money; just wanted a change in jobs; was self-employed”, I treat the separation as voluntarily induced by the worker. While all other reasons (”Company folded/changed hands/moved out of town; employer died/went out of business”, ”Strike; lockout”, ”Laid off; fired”; ”Other; transfer; any mention of armed services”, ”Job was completed; seasonal work; was a temporary job”) leads to the separation being classified as involuntary. When information on the reason why the last job ended is missing, I treat separations into unemployment as involuntary and separations into employment or out of labor force as voluntary. The question why the wife’s last job ended was asked continuously from 1979 on. This allows me to construct continuous series of quits for husbands and wives from 1978 on.

Labor earnings. The PSID reports labor income including wages and salaries, bonuses, overtime pay, tips, commissions and the like but excluding business income and farm income. To this variable, I add for each individual business income and half the household’s farm income. I deflate labor income to 1983 dollars using the CPI.

Hours worked. I use total annual work hours on all jobs including overtime as provided in the PSID (weeks worked times weekly hours plus overtime hours).

Hourly wage rate. I determine the average hourly wage rate of an individual as yearly labor earnings divided by yearly hours worked.

Contributed earnings share. I calculate the contributed earnings share as own labor earnings divided by the sum of one’s own and the partner’s labor earnings.

Year. I use year dummies to indicate years. For detrending issues, I construct a variable that runs from 1 in the first year of my main sample (1978) to 19 in the last year (1996).

Age. I use a full set of dummies for age measured in years.

Education. I use dummy variables for the following six education categories: ”less than 9 years of schooling”, ”9 - 11 grades; some high school; junior high”, ”12 grades; high school”, ”12 grades plus non-academic training or College, no degree”, ”College degree, no advanced degree mentioned”, ”College, advanced or professional degree”.

Race. I use dummies indicating white and non-white individuals, respectively.

Region. I use dummies for the four major Census regions: West (Alaska, Washington State, Oregon, Idaho, Montana, Wyoming, California, Nevada, Utah, Colorado, Arizona, New Mexico, Hawaii), Midwest (North Dakota, South Dakota, Nebraska, Kansas, Minnesota, Iowa, Missouri, Wisconsin, Illinois, Indiana, Michigan, Ohio), South (Texas, Oklahoma, Arkansas, Louisiana, Kentucky, Tennessee, Mississippi, Alabama, West Virginia, Maryland, Delaware, Washington DC, Virginia, North Carolina, South Carolina, Georgia, Florida), and Northeast (Maine, New Hampshire, Vermont, Massachusetts, Connecticut, Rhode Island, New York, Pennsylvania, New Jersey).

Children. I use dummy variables for the following 6 categories: no child, 1 child, 2 children, 3 children, 4 children, 5 or more children.

Industry. I use dummies for the following 12 major industry groups: ”Agriculture, Forestry, and Fisheries”, ”Mining”, ”Construction”, ”Manufacturing”, ”Transportation, Communications, and Other Public Utilities”, ”Wholesale and Retail Trade”, ”Finance, Insurance, and Real Estate”, ”Business and Repair Services”, ”Personal Services”, ”Entertainment and Recreation Services”, ”Professional and Related Services”, ”Public Administration”.

Occupation. I use dummies for the following 12 major industry groups: ”Professional, Technical, and Kindred Workers”, ”Managers and Administrators, except Farm”, ”Sales Workers”, ”Clerical and Kindred Workers”, ”Craftsman and Kindred Workers”, ”Operatives, except Transport”, ”Transport Equipment Operatives”, ”Laborers, except Farm”, ”Farmers and Farm Managers”, ”Farm Laborers and Farm Foremen”, ”Service Workers, except Private Household”, ”Private Household Workers”.

Tenure. I use eleven dummy variables for 0, 1, ..., 9, 10 and more completed years of tenure.

Table 13: Descriptive statistics.

	jobs held by men	jobs held by women
N	20,231	20,131
separation	0.1982	0.2439
quit	0.0844	0.1203
hourly wage rate	11.23	7.59
yearly hours worked	2261.8	1604.4
contributed earnings share	0.6695	0.3371

Notes: Hourly wage rate in 1983 dollars.

B.2 Sample selection

The sample selection is similar to the one in Bredemeier et al. (2019). I select a sample of married spouses aged 20 to 65 with male household heads. I drop the sample of economic opportunity (SEO) which is not representative for the U.S. population.

In order to handle outliers and data errors, I drop household-year observations where an individual's age falls or increases by more than two years from one year to the next, an individual's wage rate or hours worked increase by more than 250% or fall by more than 40% between two years, where an individual works more than 93 hours on average per week, or where an individual's hourly wage rate falls into the top 0.5% of the distribution.

I reshape the data to a sample of jobs with information on pay, hours, subsequent separation and the worker's socio-economic, demographic, and family background. I disregard jobs held by women whose husbands do not work in the considered year. The final sample consists of about 40,000 jobs held by married spouses. Table 13 summarizes some descriptive statistics about the final sample.

B.3 Regression specifications

Columns (1) and (2) of Table 4: Dependent variable: Quit. Independent variables: Log hourly wage rate, contributed earnings share, log hourly wage rate times year, log hourly wage rate times year squared, age dummies, year dummies, education dummies, region dummies, race dummies, children dummies. Estimated with OLS.

Columns (3) and (4) of Table 4: Dependent variable: Quit. Independent variables: Log hourly wage rate, contributed earnings share, log hourly wage rate times contributed earnings share, log hourly wage rate times year, log hourly wage rate times year squared, age dummies, year dummies, education dummies, region dummies, race dummies, children dummies. Estimated with OLS.

Row 1 of Table 5: Dependent variable of the second stage: Quit. Independent variables of the second stage: Log hourly wage rate, contributed earnings share, log hourly wage rate times contributed earnings share, log hourly wage rate times year, log hourly wage rate times year squared, age dummies, year dummies, education dummies, region dummies, race dummies, children dummies (for women only). Dependent variable of the first stage: contributed earnings share (including in its interaction with the log wage rate). Independent variables of the first stage: partner's age dummies, partner's education dummies, partner's occupation dummies, partner's race dummies, children dummies (for men only). Estimated with 2SLS.

Row 2 of Table 5: Dependent variable: Quit. Independent variables: Log hourly wage rate, contributed earnings share, log hourly wage rate times contributed earnings share, log hourly wage rate times year, log hourly wage rate times year squared, age dummies, year dummies, education dummies, region dummies, race dummies, children dummies. Estimated with OLS for a restricted sample of individuals with a positive number of children.

Row 3 of Table 5: Dependent variable: Quit. Independent variables: Log hourly wage rate, contributed earnings share, log hourly wage rate times contributed earnings share, log hourly wage rate times number of children, log hourly wage rate times year, log hourly wage rate times year squared, age dummies, year dummies, education dummies, region dummies, race dummies, children dummies. Estimated with OLS.

Row 4 of Table 5: Dependent variable: Quit. Independent variables: Log hourly wage rate, contributed earnings share, log hourly wage rate times contributed earnings share, log hourly wage rate times year, log hourly wage rate times year squared, age dummies, year dummies, education dummies, region dummies, race dummies, children dummies. Estimated with OLS for a restricted sample of individuals who remain in the labor force in year $t + 1$.

Row 4 of Table 6: Dependent variable: Quit. Independent variables: Log hourly wage rate, contributed earnings share, log hourly wage rate times contributed earnings share, log hourly wage rate times year, log hourly wage rate times year squared, age dummies, year dummies, education dummies, region dummies, race dummies, children dummies. Estimated with probit.

Row 5 of Table 6: Dependent variable: Quit. Independent variables: Log hourly wage rate, contributed earnings share, log hourly wage rate times contributed earnings share, log hourly wage rate times year, log hourly wage rate times year squared, age dummies, year dummies, education dummies, region dummies, race dummies, children dummies. Estimated with logit.

Row 6 of Table 6: Dependent variable: Quit. Independent variables: Log hourly wage rate, square of log hourly wage rate squared, cube of log hourly wage rate, contributed earnings share, log hourly wage rate times contributed earnings share, log hourly wage rate times year, log hourly wage rate times year squared, age dummies, year dummies, education dummies, region dummies, race dummies, children dummies. Estimated with OLS.

Row 7 of Table 6: Dependent variable: Quit. Independent variables: Log hourly wage rate, contributed earnings share, log hourly wage rate times contributed earnings share, log hourly wage rate times year, log hourly wage rate times year squared, age dummies, year dummies, education dummies, region dummies, race dummies, children dummies, age dummies times contributed earnings share, year dummies times contributed earnings share, education dummies times contributed earnings share, region dummies times contributed earnings share, race dummies times contributed earnings share, children dummies times contributed earnings share. Estimated with OLS.

Row 8 of Table 6: Dependent variable: Quit. Independent variables: Log hourly wage rate, contributed earnings share, log hourly wage rate times contributed earnings share, log hourly wage rate times year, log hourly wage rate times year squared, age dummies, year dummies, education dummies, region dummies, race dummies, children dummies, tenure dummies. Estimated with OLS.

Row 9 of Table 6: Dependent variable: Quit. Independent variables: Log hourly wage rate, contributed earnings share, log hourly wage rate times contributed earnings share, age dummies, year dummies, education dummies, region dummies, race dummies, children dummies. Estimated with OLS.

Row 10 of Table 6: Dependent variable: Quit. Independent variables: Log hourly wage rate, contributed earnings share, log hourly wage rate times contributed earnings share, log hourly wage rate times year, log hourly wage rate times year squared, age dummies, year dummies, education dummies, region dummies, race dummies, children dummies. Estimated with OLS. Additionally includes the year 1972 for men.

Row 1 of Table 7: Dependent variable: Quit. Independent variables: Log hourly wage rate, log own labor earnings (+1 for women), log hourly wage rate times log own labor earnings (+1 for women), log hourly wage rate times year, log hourly wage rate times year squared, age dummies, year dummies, education dummies, region dummies, race dummies, children dummies. Estimated with OLS.

Row 2 of Table 7: Dependent variable: Quit. Independent variables: Log hourly wage rate, square of log hourly wage rate, log hourly wage rate times year, log hourly wage rate times year squared, age dummies, year dummies, education dummies, region dummies, race dummies, children dummies. Estimated with OLS.

Row 3 of Table 7: Dependent variable: Quit. Independent variables: Log hourly wage rate, log yearly hours worked (+1 for women), log hourly wage rate times log yearly hours worked (+1 for women), log hourly wage rate times year, log hourly wage rate times year squared, age dummies, year dummies, education dummies, region dummies, race dummies, children dummies. Estimated with OLS.

Row 4 of Table 7: Dependent variable: Quit. Independent variables: Log hourly wage rate, partner's log labor earnings (+1 for women's earnings), log hourly wage rate times partner's log labor earnings (+1 for women's earnings), log hourly wage rate times year, log hourly wage rate times year squared, age dummies, year dummies, education dummies, region dummies, race dummies, children dummies. Estimated with OLS.

Row 5 of Table 7: Dependent variable: Quit. Independent variables: Log hourly wage rate, sample average of own and partner's labor earnings, log hourly wage rate times sum of sample average of own and partner's labor earnings in hundred thousand 1983 dollars, log hourly wage rate times year, log hourly wage rate times year squared, age dummies, year dummies, education dummies, region dummies, race dummies, children dummies. Estimated with OLS.

C Appendix to quantitative model analysis

C.1 Quit regressions in the model

While all other model evaluations only consider steady states, the quit regressions presented in Table 10 use a simulation of the model around its steady state. For this, I proceed as follows. I first solve for the steady state as described at the end of Section 4.1. This gives data for the steady-state values of wage rates, hours and marginal utility in the different cells of my model as well as cell sizes. I then simulate data as follows. I draw realizations of the idiosyncratic productivity process for each of the 20 firms, 5 in industry 1 and 15 in industry 2, for 219 periods (thus creating 19 periods of data an in my PSID analysis, after cutting 200 burn-in periods).

For each cell, I then first determine pay of every individual firm given their realized wage processes. This is simple because of the abstraction from strategic interaction and the fact that the composition of the workforce is constant within a cell (there are shocks to non-pay job preferences but, due to the law of large numbers, the distribution of these preferences is constant). I then simulate preferences for N workers where N is the share of married workers of the considered gender in this cell times the average annual sample size of my empirical data set. For each worker, I first draw initial preferences from the type-I EV distribution and then, for every period, draw workers for a re-assignment of preferences with probability θ. When drawn for re-assignment, they obtain new draws from the type-I EV distribution.

Based on simulated wage rates offered by the different firms and non-pay preferences of the sampled individuals, firm choices can be obtained as follows. I first calculate labor earnings that the worker would generate at each potential firm and translate it into utils through multiplication with the steady state marginal utility of wealth. To this I add non-pay job utility and determine the maximum of the sum. This gives time series of wage rates and chosen firms for each simulated individual. I combine the data for the different simulated individuals from the different labor-market cells to a panel data set which I then use to run quit regressions. I repeat simulation and subsequent estimation 10,000 times in a Monte-Carlo style exercise. For my baseline calibration, the results are shown in Table 10 in the main text.

Since my baseline calibration features strong autocorrelation of firm-specific productivity, I also performed Monte-Carlo estimations for alternative calibrations. In particular, I also consider

Table 14: Results of quit regressions from simulated data for alternative calibrations of the shock process for firm-specific productivity.

	men	**women**
Baseline calibration	-0.1404^{***}	-0.1481^{***}
$\rho = 0.9950$, $\sigma = 0.1200$	(0.0106)	(0.1481)
Bachmann and Bayer (2014)	-0.3404^{***}	-0.2941^{***}
$\rho = 0.9675$, $\sigma = 0.0905$	(0.0372)	(0.0368)
Bachmann et al. (2013)	-0.7428^{***}	-0.5755^{***}
$\rho = 0.8612$, $\sigma = 0.0472$	(0.0860)	(0.0859)

Mean estimates and standard errors from 10,000 Monte-Carlo repetitions. Dependent variable: quit between observation year and following year; linear probability models; standard errors in parentheses; *, **, *** indicate $p < 0.1$, $p < 0.05$, $p < 0.01$; alternative wage determinant (X): mean wage in respective cell p; additionally included variables: constant, earnings share.

the parameter values used in Bachmann and Bayer (2014) and Bachmann et al. (2013). My favorite calibration uses the values from Bachmann and Bayer (2009) since they stem from an estimation of a firm-specific productivity process while, for example, Bachmann et al. (2013) use the autocorrelation of sector-specific productivity also for firm-specific productivity. Table 14 shows the most important results, i.e. the coefficients on the interaction term between the log wage rate and the contributed share to household earnings, for the baseline and alternative calibrations. Also for lower values of the autocorrelation, the model predicts a negative coefficient here as is found in the empirical analysis in Section 3.

C.2 Additional counterfactual results

Table 15 reports the levels of gender-specific wage rates, hours worked, and elasticities of labor supply to individual firms in the different model simulations for which gender gaps are reported in Tables 11 and 12.

C.3 Extended model versions

In this appendix, I address the following issues: home production, gender differences in elasticities of labor supply to the individual firm, firm entry, and the possibility that competition for workers between firms is performed using permanent rather than period wage rates. To concentrate on the main mechanism that is active within couple households, I do so within a model without singles

Table 15: Levels of gender-specific wage rates, hours worked, and elasticities of labor supply to individual firms in the different model simulations.

	w_m	w_f	h_m	h_f	η_m	η_f
Full model (endogenous mobility)						
baseline	1.2195	1.0000	0.3466	0.3051	2.4130	1.7930
$\Delta\nu = 0$	1.2126	1.0072	0.3440	0.3184	2.3709	1.8330
$\Delta\gamma = 0$	1.2167	1.0178	0.3455	0.3064	2.3946	1.9384
$\Delta\nu = 0$ and $\Delta\gamma = 0$	1.2097	1.0250	0.3429	0.3197	2.3524	1.9811
$\Delta a \downarrow$ 10 lp	1.1982	1.1289	0.3386	0.3139	2.2847	1.9147
$\Delta\nu \downarrow$ 10 lp	1.2126	1.0073	0.3440	0.3185	2.3706	1.8333
$\Delta\gamma \downarrow$ 10 lp	1.2151	1.0281	0.3449	0.3071	2.3840	2.0276
Reference model (exogenous mobility)						
baseline	1.2195	1.0000	0.3466	0.3051	2.4130	1.7930
$\Delta\nu = 0$	1.2176	1.0016	0.3444	0.3201	2.4130	1.7930
$\Delta\gamma = 0$	1.2163	1.0720	0.3429	0.3093	2.4130	2.4130
$\Delta\nu = 0$ and $\Delta\gamma = 0$	1.2143	1.0738	0.3407	0.3244	2.4130	2.4130
$\Delta a \downarrow$ 10 lp	1.2147	1.1097	0.3410	0.3114	2.4130	1.7930
$\Delta\nu \downarrow$ 10 lp	1.2179	1.0014	0.3448	0.3176	2.4130	1.7930
$\Delta\gamma \downarrow$ 10 lp	1.2184	1.0246	0.3453	0.3066	2.4130	1.9816

and that has also no further within-gender inequality beyond non-pay job preferences. In all model versions considered here, I calibrate productivity as well as preference weights on labor supply and non-pay job attributes to match gender gaps in wage rates, hours worked, and inter-firm mobility as in the baseline evaluations, see Section 4.2, but for married individuals only. Regarding the gender gap in inter-firm mobility, I use the result of Webber (2016) that the gap is three fifth larger for married individuals.[35] Choices and calibrations for additional model parameters in the different versions are described below. In the following, I describe extended versions of the model and, thereafter, I present and compare results for the different versions.

Benchmark. For comparison, I am evaluating a model that has none of the additional aspects and no within-gender inequality. This model can be understood as the average cell for married individuals in my full quantitative model. It is described by equations (17) to (20) where the cell index m is identical to the gender index g.

Home production and gender differences in the elasticity of labor supply to the market. In this model version, households additionally produce and enjoy a home good d. The household

[35]This gives the following moments to be matched: $w_m = 1.2376$, $w_f = 0.9646$, $h_m = 0.3486$, $h_f = 0.2924$, $\eta_m = 2.6379$, and $\eta_f = 1.6402$.

target function now reads

$$u = \ln c - \delta \cdot \ln d - \frac{1}{\nu_f} \cdot \frac{(h_f + h_f^h)^{1+1/\psi}}{1+1/\psi} - \frac{1}{\nu_m} \cdot \frac{(h_m + h_m^h)^{1+1/\psi}}{1+1/\psi}$$
$$+ \frac{1}{\gamma_f} \cdot (1 - |k_f - v_f|) + \frac{1}{\gamma_m} \cdot (1 - |k_m - v_m|)\,,$$

where δ is the weight on consumption of the home good and h_g^h are hours worked in home production. The perfect substitutability of market hours and home hours follows Alesina et al. (2011). As shown by Alesina et al. (2011) this preference function endogenously gives rise to gender differences in the Frisch elasticity of labor supply to the market (rather than to the firm). Such gender differences are discussed in the empirical literature with women's labor supply *to the market* usually being found to be more elastic than men's (Keane 2011; Bredemeier et al. 2019).

The home production function is Cobb-Douglas with elasticity θ,

$$d = (h_f^h)^\theta (h_m^h)^{1-\theta}.$$

I eschew a total factor productivity level in this function as it would not be identified separately from the preference weight δ. I calibrate δ and θ to match empirical home hours by gender.[36] In this model version, I calibrate η (which is the Frisch elasticity of *total* work including housework) to maintain a Frisch elasticity of labor supply *to the market* of 0.5 for men. As an untargeted moment, the Frisch elasticity for women is about 40% larger which lies in the ballpark of gender differences in Frisch elasticities estimated by Bredemeier et al. (2019).[37]

Firm entry. Firms' net profits change in my counterfactual experiments and this may lead to changes in the number of firms. The associated changes in competition may impact on the results. In order to analyze this possibility, I alter the model as follows. I introduce a fixed cost κ which may include, among other things, supervisory labor costs as well as lump-sum fees and taxes. In every period, the number of firms V is determined by free entry total profits are zero,

$$(a_m - w_m)h_m/V - (a_f - w_f)h_f/V - \kappa = 0. \tag{24}$$

[36] Targets are $h_m = 0.1102$ and $h_f = 0.1441$ which are the average weekly hours of unpaid household work of married women and married men from the 2005 PSID expressed as shares of a weekly time endowment of 120 hours.

[37] Bredemeier et al. (2019) propose and use an estimation method that corrects for estimation biases due to borrowing constraints. They show that methods that suffer from such biases overestimate gender differences in the elasticity of labor supply to the market.

Note that V impacts on both wage rates and hours. Technically, V becomes an additional endogenous variable and (24) an additional equilibrium condition. I calibrate κ to achieve $V = 10$ which is the average number of firms per industry in the full model, see Section 4.2.

Permanent-wage competition. In the baseline model, the equilibrium concept imposes that a unilateral change in wage rates within a period does not pay off to any individual firm. One may argue that, in reality, such short-lived pay changes are hard to implement for firms independent of a potential loss of workers to other firms and that competition for workers between firms is rather performed using permanent wage rates. In this model version, I study this possibility. This implies that the elasticity of total labor supply to individual firms (with respect to the permanent wage rate) is now given by

$$\phi_p = \gamma_g \cdot V_z \cdot w_p h_p / (w_p h_p + w_{-p} h_{-p}) + \Psi_p$$

where Ψ_g is the uncompensated (Marshallian) labor-supply elasticity, instead of (18). The uncompensated elasticity is endogenously gender-specific and reads

$$\Psi_p = \eta \cdot \frac{1 - e_p}{1 + \eta e_p},$$

where e_p is the individual's contributed share to household earnings.

Comparing model versions. Table 16 summarizes the main results for the alternative model versions. For each model version, it reports the shares of the endogenous gender gaps which are created without gender differences in preferences and the shares which are to be assigned to the different dimensions of gender differences in preferences (as in Table 11 for my baseline model). It also reports the changes in gender gaps induced by counterfactually closing exogenous gender gaps by ten log points (as in Table 12 for my baseline model). As for the baseline model, I also consider a reference version where the elasticity of labor supply to individual firms is considered as exogenous.

Across the different model versions, I find my key results confirmed: First, with endogenous inter-firm mobility, the importance of gender differences in preferences is very limited as substantial shares of gender gaps in labor-market outcomes (e.g., 82-93% of the wage gap) emerge also with preference differences shut off. The importance of preference differences in the importance of non-

Table 16: Results of model versions with additional features.

	full model (endogenous mobility)			reference model (exogenous mobility)		
	Δw	Δwh	$\Delta\eta$	Δw	Δwh	$\Delta\eta$
Benchmark						
share of gap created...						
...without preference gaps	85.7%	75.4%	67.4%	57.2%	50.3%	0.0%
...by $\Delta\nu$	7.1%	18.3%	16.4%	0.0%	12.0%	0.0%
...by $\Delta\gamma$	7.3%	6.4%	16.3%	42.8%	37.7%	100.0%
change in gap induced by ...						
...10 lp reduction in Δa	−15.1 lp	−22.7 lp	−22.7 lp	−10.0 lp	−15.0 lp	0.0 lp
...10 lp reduction in $\Delta\nu$	−1.7 lp	−7.6 lp	−7.6 lp	0.0 lp	−5.0 lp	0.0 lp
...10 lp reduction in $\Delta\gamma$	−3.6 lp	−5.4 lp	−15.4 lp	−2.4 lp	−3.6 lp	−10.0 lp
Model with home production and gender gap in Frisch elasticities						
share of gap created ...						
...without preference gaps	92.6%	88.1%	78.8%	62.4%	64.6%	0.0%
...by $\Delta\nu$	1.7%	6.6%	5.9%	0.0%	5.0%	0.0%
...by $\Delta\gamma$	5.6%	5.3%	15.3%	37.6%	35.4%	100.0%
change in gap induced by ...						
...reduction in Δa	−13.4 lp	−21.5 lp	−21.5 lp	−10.0 lp	−16.0 lp	0.0 lp
...reduction in $\Delta\nu$	−0.4 lp	−2.6 lp	−2.6 lp	0.0 lp	−2.0 lp	0.0 lp
...reduction in $\Delta\gamma$	−2.8 lp	−4.5 lp	−14.5 lp	−2.1 lp	−3.3 lp	−10.0 lp
Model with firm entry						
share of gap created ...						
...without preference gaps	85.7%	75.4%	67.4%	57.2%	50.3%	0.0%
...by $\Delta\nu$	7.3%	8.5%	6.5%	0.0%	12.0%	0.0%
...by $\Delta\gamma$	7.2%	6.3%	16.2%	42.8%	37.7%	100.0%
change in gap induced by ...						
...10 lp reduction in Δa	−15.2 lp	−22.8 lp	−22.8 lp	−10.0 lp	−15.0 lp	0.0 lp
...10 lp reduction in $\Delta\nu$	−1.8 lp	−7.7 lp	−7.7 lp	0.0 lp	−5.0 lp	0.0 lp
...10 lp reduction in $\Delta\gamma$	−3.5 lp	−5.3 lp	−15.3 lp	−2.4 lp	−3.6 lp	−10.0 lp
Model with permanent-wage competition						
share of gap created ...						
...without preference gaps	82.7%	72.8%	65.1%	49.0%	43.1%	0.0%
...by $\Delta\nu$	8.0%	19.1%	17.1%	0.0%	12.1%	0.0%
...by $\Delta\gamma$	9.4%	8.3%	18.0%	51.0%	44.9%	100.0%
change in gap induced by ...						
...10 lp reduction in Δa	−15.8 lp	−23.8 lp	−23.8 lp	−10.0 lp	−15.0 lp	0.0 lp
...10 lp reduction in $\Delta\nu$	−2.0 lp	−7.9 lp	−7.9 lp	0.0 lp	−5.0 lp	0.0 lp
...10 lp reduction in $\Delta\gamma$	−4.6 lp	−6.9 lp	−16.9 lp	−2.9 lp	−4.4 lp	−10.0 lp

Notes: Models calibrated to married couples. In models with home production and with permanent-wage competition, reference model treats the elasticity of total labor supply to individual firm (including intensive margin) as exogenous. $\Delta x = \ln x_m - \ln x_f$. lp = log points.

pay job attributes for the gender gap in inter-firm mobility is only between 15% and 18%. Models that mistake inter-firm mobility as exogenous, by contrast, strongly overestimate the importance of preference differences. The reference models with exogenous mobility, e.g., assign only 49-63% of the gender wage gap to non-preference factors. They particularly overestimate the role of gender differences in the importance of non-pay job attributes to which they assign 37-51% of the gender wage gap and, by construction, 100% of the gender gap in inter-firm mobility. Second, the relation of inter-firm mobility to earnings positions within the household leads to substantial amplification across model versions. Changes in exogenous labor-demand factors (a) or exogenous mobility factors (γ) have effects on the gender wage gap and the gender gap in inter-firm mobility which are 33-59% stronger in the models with endogenous mobility compared to the models with exogenous mobility. There are also moderate effects of changes in exogenous labor-supply factors (ν) on the gender gaps in wage rates and inter-firm mobility which are completely overlooked by the models which mistake inter-firm mobility as exogenous.