

The Gender Pay Gap: Micro Sources and Macro Consequences

Iacopo Morchio
Christian Moser

Discussion Paper 21/751

14 March 2021

School of Economics

University of Bristol
Priory Road Complex
Bristol
BS8 1TU
United Kingdom



The Gender Pay Gap: Micro Sources and Macro Consequences*

Iacopo Morchio[†]

Christian Moser[‡]

March 14, 2021

Abstract

We document that a large share of the gender pay gap in Brazil is due to women working at lower-paying employers. However, compared with that of men, women's revealed-preference ranking of employers is less increasing in pay. To interpret these facts, we develop an empirical equilibrium search model with endogenous gender differences in pay, amenities, and recruiting intensities across employers. The estimated model suggests that compensating differentials explain one-fifth of the gender pay gap, that there are significant output and welfare gains from eliminating gender differences, and that equal-treatment policies fail to close the gender pay gap.

Keywords: Empirical Equilibrium Search Model, Linked Employer-Employee Data, Worker and Firm Heterogeneity, Misallocation, Compensating Differentials, Discrimination

JEL Classification: E24, E25, J16, J31

*We thank Jim Albrecht, Jesper Bagger, Erling Barth, Mark Colas, Rafael Dix-Carneiro, Carlos Carrillo-Tudela, David Dorn, Jean Flemming, Nezih Guner, Jonathan Heathcote, Kyle Herkenhoff, Francis Kramarz, Kory Kroft, Ilse Lindenlaub, Paolo Martellini, Guido Menzio, Andreas Mueller, Makoto Nakajima, Josep Pijoan-Mas, Laura Pilossoph, Tommaso Porzio, Todd Schoellman, Kjetil Storesletten, Coen Teulings, Gabriel Ulyssea, Diana Van Patten, Chris Taber, Rune Vejlin, Gianluca Violante, and Jonathan Vogel for valuable comments and discussions. We also thank attendants at the 2018 GEA Conference, the 2019 AEA Annual Meeting, the 2019 Search and Matching Network Conference, the 2019 Bank of Italy/CEPR/IZA Conference, the 2019 SED Annual Meeting, the 2019 and 2020 NBER Summer Institutes, the 2019 and 2020 Stanford SITE Workshops, the 3rd Dale T. Mortensen Centre Conference, the Boston Macro Juniors Conference, the 2020 Columbia Junior Micro Macro Labor Conference, the EALE/SOLE/AASLE World Conference in Berlin, the 5th Annual USC Marshall Macro Day, and seminar participants at Columbia University, Boston University, CREST, Iowa State University, the University of Oslo, the University of Bristol, the University of Edinburgh, the University of Vienna, the Johannes Kepler University Linz, the Federal Reserve Bank of Philadelphia, Temple University, the Federal Reserve Bank of Minneapolis, the VMACS Junior Conference, the University of Western Ontario, Arizona State University, Emory University, the Federal Reserve Bank of Atlanta, the Essex/RHUL/Bristol SaM Webinar, the Federal Reserve Bank of St. Louis, Banco de Portugal, CEMFI, University of Bonn, and UCL for helpful comments. Ian Ho and Rachel Williams provided outstanding research assistance. Moser gratefully acknowledges financial support from the Ewing Marion Kauffman Foundation and the Sanford C. Bernstein & Co. Center for Leadership and Ethics at Columbia Business School. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System. Any errors are our own.

[†]University of Bristol. Email: iacopo.morchio@bristol.ac.uk.

[‡]Columbia University and Federal Reserve Bank of Minneapolis. Email: c.moser@columbia.edu.

1 Introduction

During the past decades, the introduction of gender in economic theory and measurement has had a profound impact on studies of labor markets and the macroeconomy. A common thread in these studies is the robust empirical finding of a gender pay gap, which is explained partly by gender imbalances in employment across different types of jobs. The implications of this finding depend crucially on whether the observed pay and employment of women relative to men reflects gender-specific preferences over jobs or barriers in the labor market. The goal of this paper is to identify the microeconomic sources of the gender pay gap and to assess its macroeconomic consequences.

Our contribution is threefold. First, we establish novel facts on employment segregation, pay heterogeneity, and revealed-preference ranks by gender across employers. Second, we interpret these facts by developing and estimating a new empirical equilibrium search model featuring endogenous gender differences in pay, amenities, and recruiting intensities across employers. Third, we use the estimated model to decompose the empirical gender pay gap, to quantify the output and welfare gains from moving to an economy with no gender differences, and to evaluate the effects of counterfactual equal-treatment policies, including equal-pay, equal-hiring, and equal-amenities mandates. In doing so, we provide the first estimates of output and welfare losses from firm-level gender misallocation.

We base our empirical analysis on rich linked employer-employee data from Brazil between 2007 and 2014. Brazil's large gender earnings gap of around 14 log points makes it interesting in its own right to study the sources of gender inequality in a nation of over 200 million people. Such a study is made possible by Brazil's remarkable data infrastructure, which contains detailed information on gender-relevant labor market variables, including workers' education, occupation, tenure, work hours, and employment histories with information on parental leaves.

To dissect the gender pay gap, we first estimate a regression with gender-specific employer pay components developed by [Card et al. \(2016\)](#), who in turn build on the seminal two-way fixed effects (FEs) framework by [Abowd, Kramarz, and Margolis \(1999\)](#), henceforth AKM). Controlling for worker heterogeneity, we find a gender pay gap of around 8 log points, which is accounted for by gender-specific employer pay heterogeneity, with women sorting to lower-paying employers relative to men.

Our next objective is to assess whether the gender pay gap across employers reflects gender-specific tastes or obstacles. To this end, we construct revealed-preference rankings of employers using PageRanks ([Page et al., 1998](#); [Sorkin, 2018](#)) estimated separately by gender. The PageRank is a network centrality measure that quantifies the attractiveness of an employer based on the network of

worker flows between employers. Intuitively, higher-ranked employers poach many workers from other high-ranked employers and lose few workers to low-ranked employers. Importantly, these PageRanks estimates are independent of any information on employer sizes or pay.

Based on estimated employer pay and revealed-preference ranks, we establish three novel facts. First, for both genders, employment is concentrated among high-ranked, but not necessarily high-paying, employers. That is, employers' revealed-preference ranks differ from pay ranks for both women and men. Second, revealed-preference employer ranks are increasing in pay for both genders, but more so for men than for women. In other words, women are relatively more attracted to employers with low pay but high values of nonpay characteristics. Third, men and women agree more on pay and revealed-preference rankings toward the top, compared with further down the employer rank distribution. Together, these facts shed light on the microstructure of gender-specific pay and nonpay attributes across employers.

To interpret these facts, we develop a new empirical equilibrium search model featuring endogenous gender differences in pay, amenities, and recruiting intensities across employers. The model accommodates several competing explanations for the gender pay gap—including employer productivity differences (Burdett and Mortensen, 1998), gender-specific compensating differentials (Rosen, 1986), statistical discrimination (Arrow, 1971) based on expected employment transitions, and taste-based discrimination (Becker, 1971)—while remaining analytically tractable. Our framework gives rise to gender-specific job ladders with several notable properties. The equilibrium wage equation is log-additively separable in a worker component and a gender-specific employer component, thereby providing a microfoundation for our reduced-form specification. Endogenous worker transitions may be associated with wage declines. Gender discrimination can survive in the presence of labor market frictions. Finally, general-equilibrium forces imply that even employers without regard for gender offer different pay to men and women.

A prerequisite for assessing the macro consequences of gender inequality is the separate identification of competing explanations for the observed gender gaps. To this end, we develop a novel estimation method that combines measures of employer pay and revealed-preference ranks to estimate over three million parameters guiding employer and worker heterogeneity in four categories. First, we deduce gender-specific employer amenity values as residuals between employers' relative pay and revealed-preference ranks. Second, gender-specific labor market parameters are identified from empirical worker flows across employer ranks. Third, employers' preferences over gender are identified from the conditional equilibrium pay gap between men and women within establishments.

Finally, we estimate gender-specific employer recruiting costs based on empirical hiring patterns.

The estimated model is qualitatively and quantitatively consistent with our empirical facts. By flexibly estimating the model, we shed light on unobserved worker and employer heterogeneity, without further restrictions on functional forms. This allows us to uncover that employer pay, ranks, and amenities are positively correlated within employers across genders. Employer ranks are relatively more increasing in pay for men but more increasing in amenity values for women. We find compensating differentials for both genders. Employers' preference for men over women increases with employer productivity, consistent with [Becker's \(1971\)](#) idea that discrimination cannot survive among low-productivity firms with close-to-zero profits.

Our novel firm-level identification allows us to link structural estimates of employer heterogeneity to observed employer characteristics in the data. We find that women put relatively higher value on amenities like hours flexibility and parental leave benefits, while men are relatively less averse to pay fluctuations and health risks. Larger employers face relatively lower costs of recruiting women, while those more likely to rely on referral networks and those based in commute-heavy environments are more biased toward hiring men. Empirical proxies for more woman friendly employers include higher routine manual and nonroutine cognitive-interpersonal task intensities, higher female employment shares, employing women in top-paid positions, and greater financial accountability. These estimates speak to different reasons why some employers are not gender-blind, including taste-based discrimination ([Becker, 1971](#)) and gender-specific comparative advantages ([Goldin, 1992](#)).

With the estimated equilibrium model in hand, we simulate a number of counterfactuals that shed light on the sources and consequences of the gender pay gap. In a nonlinear structural decomposition, we find that compensating differentials in the form of gender-specific amenities explain 1.3 log points (18 percent) of the gap, employer tastes explain 5.4 log points (73 percent), and gender-specific hiring costs account for 5.6 log points (76 percent) of the gap. However, given the estimated distribution of pay and nonpay characteristics across employers, closing the gender pay gap may or may not be welfare improving. We find that moving to an economy without gender differences is associated with significant output gains of 3.5 percent and welfare gains of 3.3 percent. In contrast, hypothetical equal-treatment policies—equal-pay, equal-hiring, and equal-amenities mandates—yield mild and at times negative consequences on average, though these policies have important redistributive effects in equilibrium. Their ineffectiveness is explained largely by the offsetting behavioral responses of employers and the resulting changes in worker sorting across genders. Thus, our results underline the importance of studying such policies in general equilibrium.

Related literature. Several macroeconomic studies have focused on the drivers of trends in female labor force participation, including structural change (Ngai and Petrongolo, 2017), culture (Fernández et al., 2004), technology (Albanesi and Olivetti, 2016), and information (Fogli and Veldkamp, 2011). Related studies have linked changes in female participation to economic growth (Heathcote et al., 2017), unemployment (Albanesi and Şahin, 2018), business cycles (Albanesi, 2020; Fukui et al., 2020), and declining dynamism (Peters and Walsh, 2019). Whereas previous work has focused on data at the national, geographic, sectoral, or occupational level, our analysis sheds light on much more granular firm-level determinants of gender pay differences and their macro consequences.

The firm is a natural unit of analysis for studying factor input misallocation in relation to macroeconomic outcomes (Restuccia and Rogerson, 2008; Hsieh and Klenow, 2009; Lentz and Mortensen, 2008; Bagger et al., 2014). If gender-specific barriers impede women’s relocation to higher-productivity firms, then the gender pay gap may be associated with efficiency losses from misallocation of talent (Hsieh et al., 2019). Yet pay gaps or output losses do not necessarily reflect welfare losses if there are compensating differentials. By developing a new equilibrium model and estimating it on detailed linked employer-employee data, we provide the first estimates of output and welfare losses from gender misallocation, which depend crucially on the firm-level determinants of the gender pay gap.

A burgeoning literature based on AKM’s contribution highlights firm heterogeneity in explaining empirical worker pay dispersion (Card et al., 2013; Alvarez et al., 2018; Song et al., 2018). We build on Card et al.’s (2016) variant of this framework with gender-specific employer pay components to decompose the employer pay gap into sorting and rent sharing terms. By microfounding this specification based on worker and firm optimization, our equilibrium model rationalizes gender-specific sorting and rent-sharing patterns and hence the empirical gender pay gap.

Our empirical equilibrium search model builds on the pioneering framework of Burdett and Mortensen (1998). Bontemps et al. (1999, 2000) estimate variants of this framework. Other extensions and empirical applications include Moscarini and Postel-Vinay (2013), Meghir et al. (2015), Engbom and Moser (2018), and Bagger and Lentz (2019). In all these models, firms are ex ante heterogeneous in only one dimension: productivity. Consequently, all workers agree on a pay-based ranking of firms. To interpret our empirical facts on heterogeneity in employer pay and ranks by gender, we develop and estimate a tractable model with multiple dimensions of unrestricted firm heterogeneity. We find that this richer model has radically different implications for the sources and consequences of the gender pay gap than those of a model with only productivity heterogeneity.

Other models have addressed gender issues in the labor market. For example, Black (1995),

Bowlus (1997), Albanesi and Olivetti (2009), Flabbi (2010), and Amano-Patiño et al. (2019) study different forms of wage discrimination. It is well known that discrimination is hard to empirically distinguish from unobserved productivity differences or compensating differentials. While our goal is not to pinpoint discrimination, we nonparametrically identify the entire distribution—not just population parameters—of employer preferences over gender through the lens of our structural model. We argue that these preferences may reflect gender discrimination or comparative advantages and provide empirical evidence for and against both channels.

Goldin (2014) and Erosa et al. (2019) have empirically studied gender-specific compensating differentials à la Rosen (1986). They highlight the gender-specific value of job flexibility, which we confirm in our structural estimation. Hwang et al. (1998), Lang and Majumdar (2004), and Albrecht et al. (2018) develop theoretical models with *nonspecific* job amenities. Using survey data, Sullivan and To (2014), Jung and Kuhn (2017), Hall and Mueller (2018), and Luo and Mongey (2019) estimate nonspecific amenity values. These studies are silent on specific job characteristics underlying these amenity values. Others have estimated the values of *specific* job amenities like employer health insurance (Dey and Flinn, 2005), job security (Jarosch, 2015), fatality risk (Lavetti and Schmutte, 2018), commuting costs (Flemming, 2020), region (Heise and Porzio, 2019), sexual harassment (Folke and Rickne, 2020), and working conditions (Bonhomme and Jolivet, 2009). These studies do not relate the value of specific amenities to the overall value of job amenities. A strength of our firm-level identification strategy is that it allows us to bridge these two strands of the literature by first estimating nonspecific amenity values firm by firm and then relating them back to specific employer characteristics.

In related work, Xiao (2020) and Taber and Vejlín (2020) use linked employer-employee data to estimate amenity values. Our work takes a step forward by identifying gender-specific employer heterogeneity without relying on distributional assumptions or indirect inference. Uncovering the entire distribution of job amenities and other dimensions of employer heterogeneity is crucial to accurately assess the macro consequences of the gender pay gap. Sorkin (2018) and Lamadon et al. (2019) also estimate firm-level amenity values. One of our contributions is that we unpack the amenity black box by relating our structural estimates of amenity values to empirical proxies for employer amenities. Sorkin (2017) combines PageRanks and a partial-equilibrium framework to highlight gender differences in the exogenous offer distribution of pay and amenities. Such a partial-equilibrium framework is a useful accounting tool but leaves open many interesting macroeconomic questions. Complementary prior work, we leverage an equilibrium model of endogenous gender differences in the distributions of pay, amenities, and vacancies to study the effects of counterfactual policies.

Outline. The paper is structured as follows. Section 2 introduces the data. Section 3 presents empirical facts. Section 4 develops an equilibrium model. Section 5 outlines the identification strategy. Section 6 shows estimation results. Section 7 conducts counterfactuals. Finally, Section 8 concludes.

2 Data

2.1 Dataset and Variables

Our main data source is the *Relação Anual de Informações Sociais's* (RAIS) linked employer-employee register of all tax-registered firms, which is administered by the Brazilian Ministry of Labor and Employment. The data are available from 1985 onward. Since 2007, the data contain detailed information on reasons and lengths of worker absences, including parental leaves. In 2015, the country entered a severe recession. Therefore, we focus on the eight-year period from 2007 to 2014.

The data contain identifiers for workers and establishments.¹ For each job spell in each calendar year, we observe the start and end dates, mean monthly earnings (henceforth “earnings”), and contractual work hours (henceforth “hours”). We construct hourly wages (henceforth “wages”) as earnings divided by hours. Other key variables include gender, race in five categories, nationality in 37 categories, educational attainment in nine categories, worker age in years, five-digit sector codes with 672 categories, municipality codes with 5,565 categories, six-digit occupation codes with 2,383 categories, and tenure in years. We exploit the full panel dimensions of the data going back to 1985 with the tenure variable to impute actual—not just potential—formal-sector work experience in years.²

2.2 Sample Selection

We restrict attention to male and female workers between the ages of 18 and 54 who worked at least one hour per week with earnings at or above the federal minimum wage. We keep for each worker-year combination the highest paid among all longest employment spells. We then impose additional selection criteria that improve the statistical properties of our reduced-form model. We iteratively drop singleton observations defined by gender-establishment combinations and worker identifiers. We also impose a minimum establishment size threshold of 10 nonsingleton workers³ per year on

¹All of our analysis is at the level of the establishment, which we refer interchangeably to as “employer” or “firm.”

²The distinction between actual and potential experience is important given Brazil’s sizable informal sector, as shown in Figure A.1 in Appendix A.1, but explains little of the empirical gender pay gap.

³Following Sorkin (2018), nonsingleton workers are those who are observed at least one more time at a future date.

average.⁴ Finally, we require that establishments appear in our sample in at least four out of the eight years. Together, these selection criteria ensure that we are dealing with a set of reasonably large and stable establishments for which pay policies and employer ranks can be credibly estimated while minimizing limited-mobility bias (Andrews et al., 2008, 2012; Bonhomme et al., 2020). Finally, to separately identify worker and employer pay components and rank employers, we focus on observations in the largest set of strongly connected observations; doing so requires worker flows both into and out of all establishments in the connected set. We confirm that our selection criteria do not substantially change the empirical gender pay gap.

2.3 Summary Statistics

Table 1 presents summary statistics on the connected set from 2007 to 2014.⁵ The pooled sample comprises over 231 million worker-years, corresponding to around 55 million unique workers and 222 thousand unique establishments. Around 38 percent of these observations are for women. The gender gap in raw earnings is around 14 log points, and that in wages is around 6 log points. Compared with women, men are more likely to be nonwhite, have fewer years of schooling, are significantly younger, are employed at smaller establishments, work more hours, and have lower tenure.

3 Employer Heterogeneity and the Gender Pay Gap

The goal of this section is to juxtapose gender-specific pay and revealed-preference ranks across employers. A classical Mincerian analysis of the gender gap in pay is presented in Appendix B.1. Worker observable controls explain only around one-third of the empirical gender gap of over 14 log points. At the same time, Appendix B.2 shows that women in Brazil are highly segregated across employers, even within industries and occupations. For example, for over 28 percent of (large) establishments, women constitute employ less than 10 percent of the workforce. This leads us to investigate whether there is a connection between employer heterogeneity and the gender pay gap.

⁴While the RAIS data cover only Brazil's formal sector, the employer size restriction implies that the vast majority of informal establishments would in any case be excluded from our analysis (Ulyssea, 2018; Dix-Carneiro et al., 2019).

⁵Appendix A.2 compares summary statistics over the sample period (Table A.1), those before sample selections and restriction to the connected set (Table A.2), and raw data with the selected sample in the connected set (Table A.3).

Table 1. Summary statistics, 2007–2014

	Overall	Men	Women
Mean log real monthly earnings (std. dev.)	7.237 (0.709)	7.291 (0.712)	7.150 (0.695)
Mean log real hourly wage (std. dev.)	3.525 (0.758)	3.548 (0.747)	3.487 (0.774)
Share nonwhite	0.384	0.416	0.333
Mean years of education (std. dev.)	11.0 (3.3)	10.4 (3.4)	12.1 (3.0)
Mean age (std. dev.)	33.9 (9.5)	33.6 (9.5)	34.2 (9.5)
Mean establishment size (std. dev.)	3,135 (17,720)	1,860 (12,160)	5,216 (24,046)
Mean contractual work hours (std. dev.)	41.5 (5.4)	42.5 (4.0)	39.8 (6.7)
Mean years of tenure (std. dev.)	4.1 (5.9)	3.7 (5.5)	4.8 (6.4)
Mean log gender earnings gap	0.141		
Mean log gender wage gap	0.062		
Number of worker-years	231,805,831	143,745,869	88,059,962
Number of unique workers	55,078,455	33,197,634	21,880,821
Number of unique establishments	222,695	153,081	69,614
Share female	0.380		

Note: Table reports summary statistics for the pooled sample from 2007 to 2014 separately for the overall population, for men only, and for women only. Since information on race is missing for a significant share of observations, conditional means are reported for race.

3.1 Measuring Gender-Specific Employer Pay

To understand the link between employer segregation and the gender pay gap, we estimate a wage equation with gender-specific employer pay components developed by [Card et al. \(2016\)](#), building on the two-way FEs specification by AKM. This approach allows for the possibility that a given employer has two pay policies—one for each gender. Formally, we model earnings of individual i in year t working at establishment $j = J(i, t)$, denoted y_{ijt} , as

$$y_{ijt} = X_{it}\beta + \alpha_i + \mathbf{1}[\text{gender}_i = M] \psi_j^M + \mathbf{1}[\text{gender}_i = F] \psi_j^F + \varepsilon_{ijt}, \quad (1)$$

where X_{it} is a vector of gender-specific worker characteristics including a set of restricted education-age dummies as well as dummies for hours, occupation, tenure, actual experience, and education-year combinations;⁶ α_i is a person FE; ψ_j^M and ψ_j^F are the male and female employer FEs, respectively; and ε_{ijt} is a residual term.⁷ By including person FEs, we control for selection of men and women across establishments based on unobserved time-invariant worker characteristics such as ability. In estimating equation (1), our focus lies in the distribution of gender-specific employer FEs ψ_j^M and ψ_j^F .

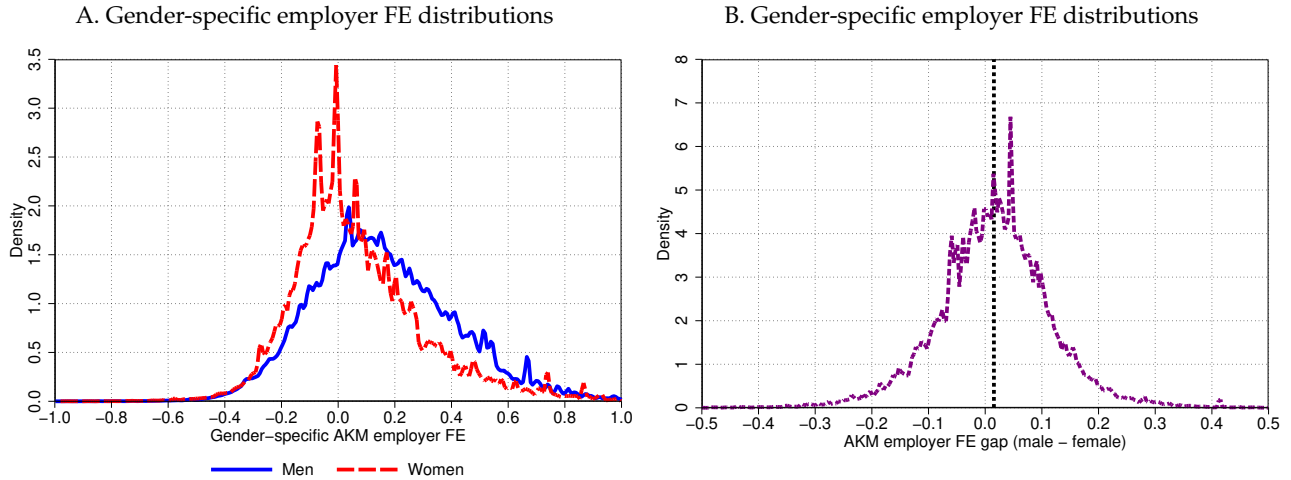
⁶By controlling for occupation, among other covariates, we abstract from the share of the gender pay gap that is due to occupational sorting in the presence of gender-specific barriers to entry.

⁷To simultaneously identify age, time, and worker FEs, we restrict the age-pay profile be flat around ages 45–49. This is an attractive alternative to the methods proposed in [Card et al. \(2018\)](#), which allows us to verify that our restriction leads to smooth education-age FEs around this age window. See [Figure B.12](#) in [Appendix B.4](#) for details.

As is the case in all two-way fixed models, at least one normalization must be made regarding the intercept or mean of the employer FEs versus the person FEs. In our case, the model with gender-specific employer FEs requires two normalizations—one for each gender. Consistent with the theoretical model presented later, we follow [Card et al. \(2016\)](#) and [Gerard et al. \(2018\)](#) in normalizing the employer FEs of both genders to be of mean zero in the restaurant and fast-food sector, which arguably is populated by low-surplus employers.⁸

We now turn to our object of interest in equation (1)—namely, the gender-specific employer FEs.⁹ Panel A of Figure 1 plots the distribution of employer FEs by gender. The distribution for women has visibly lower mean and lower variance than that for men. Panel B of the figure shows the distribution of within-employer differences in FEs for dual-gender establishments. The distribution is relatively dispersed compared with its mean of around 2 log points. Altogether, this evidence suggests that sorting of women into lower-paying employers is a significant source of gender pay differences.

Figure 1. Predicted AKM employer FEs for women and men



Note: Figure shows kernel density plots of estimated gender-specific employer FEs based on estimating earnings equation (1). Dashed vertical line shows mean of the distribution. See Figure B.16 for the different, but related, gender-specific distributions of employment across AKM employer FEs of a fixed gender.

To dissect the structure of pay and the relative contribution of employer heterogeneity, Table 2 presents a variance decomposition of log earnings.¹⁰ Men have a slightly higher variance of earnings,

⁸For robustness, we experimented with alternative normalizations for gender-specific employer FEs. Separately, we have repeated our analysis based on a wage equation without worker FEs, making the normalization redundant.

⁹In Appendix B.4, we present auxiliary results relating to the AKM equation, including estimated gender-specific hours FEs (Figure B.7), occupation FEs (Figure B.8), actual-experience FEs (Figure B.9), tenure FEs (Figure B.10), education-year FEs (Figure B.11), and education-age FEs (Figure B.12). Further tests of the log-additivity and exogenous mobility assumptions of similar specifications and data are presented in [Alvarez et al. \(2018\)](#) and [Gerard et al. \(2018\)](#).

¹⁰Table 2 shows plug-in estimators of the variance components. In ongoing work, we implement the leave-one-out estimator by [Kline et al. \(2019\)](#), which uses a jackknife bias correction for limited-mobility bias as an alternative to the methods

with 52.4 log points, compared with 51.1 log points for women. For both genders, the largest variance component is due to estimated worker FEs, which account for 24 percent for men and 25 percent for women. Employer FEs account for 12 percent of the variance of earnings for men and 11 percent of that for women. The correlation between person and employer FEs is around 23 percent for men and 27 percent for women. For each gender, the largest connected set spans close to the full data. Finally, the model explains around 93 percent of the variation in log earnings.

Table 2. Variance decomposition based on gender-specific employer FEs model

Variances	Men		Women	
	Level	Share	Level	Share
Variance of log earnings	0.524		0.511	
Components of variance of log earnings:				
Person FEs	0.125	0.238	0.129	0.252
Employer FEs	0.064	0.122	0.056	0.109
Education-year FEs	0.002	0.004	0.001	0.001
Education-age FEs	0.061	0.116	0.055	0.107
Hours FEs	0.000	0.001	0.001	0.002
Occupation FEs	0.013	0.025	0.012	0.023
Tenure FEs	0.003	0.006	0.005	0.010
Actual experience FEs	0.034	0.065	0.023	0.045
Covariances	0.185	0.352	0.198	0.387
Residual	0.037	0.071	0.033	0.064
Correlation person/employer FEs	0.226		0.270	
Observations	143,745,869		88,059,962	
R^2	0.929		0.936	

Note: Table shows decomposition of the variance of earnings based on earnings equation (1).

From here on, we focus on differences in the gender-specific employer components (henceforth “gender gap”). Using a Kitagawa-Oaxaca-Blinder decomposition, we can write the gender gap as

$$\begin{aligned}
\gamma^e &\equiv \mathbb{E}_{i,t} \left[\psi_{J(i,t)}^M \mid gender_i = M \right] - \mathbb{E}_{i,t} \left[\psi_{J(i,t)}^F \mid gender_i = F \right] \\
&= \underbrace{\mathbb{E}_{i,t} \left[\psi_{J(i,t)}^M - \psi_{J(i,t)}^F \mid gender_i = M \right]}_{\text{within-employer gender pay gap}} + \underbrace{\left(\mathbb{E}_{i,t} \left[\psi_{J(i,t)}^F \mid gender_i = M \right] - \mathbb{E}_{i,t} \left[\psi_{J(i,t)}^F \mid gender_i = F \right] \right)}_{\text{between-employer gender pay gap}} \quad (2)
\end{aligned}$$

$$\begin{aligned}
&= \underbrace{\mathbb{E}_{i,t} \left[\psi_{J(i,t)}^M - \psi_{J(i,t)}^F \mid gender_i = F \right]}_{\text{within-employer gender pay gap}} + \underbrace{\left(\mathbb{E}_{i,t} \left[\psi_{J(i,t)}^M \mid gender_i = M \right] - \mathbb{E}_{i,t} \left[\psi_{J(i,t)}^M \mid gender_i = F \right] \right)}_{\text{between-employer gender pay gap}}. \quad (3)
\end{aligned}$$

Equations (2) and (3) are alternative decompositions of the gender pay gap, γ^e , into two terms. The

proposed in [Bonhomme et al. \(2019\)](#) and [Borovičková and Shimer \(2020\)](#). Reassuringly, Appendix B.5 demonstrates that our results are stable for a range of sample restrictions designed to ease the threat of limited-mobility bias, consistent with our estimates of positive correlations between estimated person and gender-specific employer FEs of 0.226 for men and 0.270 for women.

within-employer pay gap or *pay-policy component* is the mean difference in gender-specific employer FEs weighted by the distribution of men or women. It reflects differences in pay between women and men at the same establishment. The *between-employer pay gap* or *sorting component* is the gender-weighted difference in mean male-employer FEs or female-employer FEs. It reflects differences in pay between men and women that are due to their different allocations across establishments.¹¹

Results of the two decompositions in equations (2) and (3) are shown in Table 3.¹² Out of the total gender pay gap of 8.4 log points, 24 percent (5 percent) is attributed to the pay-policy component in Decomposition 1 (2). The remainder is attributed to the sorting component. This evidence suggests that compared with men, women systematically work at lower-paying employers.¹³

Table 3. Kitagawa-Oaxaca-Blinder decompositions of the gender employer pay gap

	Gender pay gap	Pay-policy component		Sorting component	
		Level	Share	Level	Share
Decomposition 1	0.084	0.020	0.241	0.064	0.759
Decomposition 2	0.084	0.004	0.047	0.080	0.953

Note: Table shows Kitagawa-Oaxaca-Blinder decomposition of overall gender pay gap into pay-policy and sorting components. Decomposition 1 corresponds to equation (2) and uses women’s employer FEs for computing the between-employer component. Decomposition 2 corresponds to equation (3) and uses men’s employer FEs for computing the between-employer component.

3.2 Measuring Gender-Specific Revealed-Preference Employer Ranks

To what extent does the gender pay gap reflect a gender utility gap? To answer this question, one must take into account both pay and nonpay characteristics of jobs for both genders. To this end, we use the PageRank index to estimate gender-specific revealed-preference rankings of employers. The PageRank is a network centrality measure developed by Page et al. (1998) to rank websites for the web search engine Google. PageRank was first used in an economic context by Sorkin (2018).

Let $g \in \{M, F\}$ index a worker’s gender, let $j \in \mathcal{J}^g = \{j_1, j_2, \dots, j_{N^g}\}$ index a set of N^g gender-specific employers, and let $t \in \mathcal{T}$ index time. Let $n_{j,j',t}^g$ denote the number of workers of gender g transitioning from employer j to employer j' at time t , $n_{j,j'}^g = \sum_{t \in \mathcal{T}} n_{j,j',t}^g$ denote the time aggregation

¹¹The sorting component is invariant to the choice of the normalization of gender-specific employer FEs. Coincidentally, this will be the main object of interest in our study. The pay-policy component, on the other hand, depends on the normalization of men’s employer FEs relative to those of women, as discussed above.

¹²Figure B.16 in Appendix B.6 illustrates estimates of the two components of the two decompositions.

¹³To address the role of childbirth and maternity, in Appendix B.7, we study life-cycle patterns of employer pay by gender and parental status. In Appendix B.8, we conduct an event study analysis around childbirth, following the methodology by Kleven et al. (2016). While we find significant gender gaps associated with childbirth in participation and earnings, our analysis suggests that firm pay heterogeneity is *not* a very important factor behind these gaps. While it would be interesting in its own right to explicitly study these patterns, we will implicitly account for them in our structural analysis.

of gender-specific flows between the two employers, and $n_{j'}^g = \sum_j n_{j,j'}^g$ denote the number of workers of gender g flowing out from employer j . Let $B^g(j) = \{j' : n_{j',j}^g \geq 1\}$ denote the set of employers who have ever lost a worker of gender g to employer j . Let $d \in [0, 1]$ be a damping factor. The *PageRank index*, $s^g(j)$, is a probability distribution over all employers $j \in \mathcal{J}^g$ such that

$$s^g(j) = \frac{1-d}{N^g} + d \sum_{j' \in B^g(j)} w_{j',j}^g s^g(j'), \quad \forall j \in \mathcal{J}^g, \forall g, \quad (4)$$

where $w_{j',j}^g = n_{j',j}^g / n_{j'}^g$ is a weight equal to the share of worker flows from employer j' to employer j as a fraction of all worker flows from employer j' , excluding transitions between employers with intermittent nonemployment spells. Intuitively, employers with a high PageRank index poach many workers from other employers with high PageRank indices and lose few workers to employers with low PageRank indices. The damping factor d represents the weight on the poaching term in a convex combination with equal employer weights. Based on PageRank indices, we compute gender-specific *PageRanks* $r^g(j)$ for every employer $j \in \mathcal{J}^g$ as the rank of the PageRank indices, with the lowest rank normalized to 0 and the highest rank normalized to 100.

Following [Sorkin \(2018\)](#), we choose as the damping factor $d = 1$. Under this parameterization, the PageRank index equals the asymptotic share of time a representative worker (“random surfer”) who switches jobs by following the network of empirical worker flows would spend at a given employer. This interpretation of PageRanks matches the definition of employer ranks in a large class of on-the-job search models, including the one we develop. Note also that an employer’s PageRank does not mechanically depend on its pay or size. Indeed, in computing PageRanks, we did not use any information on worker wages or the number of workers at any employer. Based on equation (4), we compute unweighted employer PageRanks separately by gender.¹⁴

3.3 Linking Gender-Specific Employer Pay to Revealed-Preference Employer Ranks

In this section, we combine estimates of gender-employer FEs in pay and PageRanks to establish three facts on employer heterogeneity in pay and ranks.¹⁵

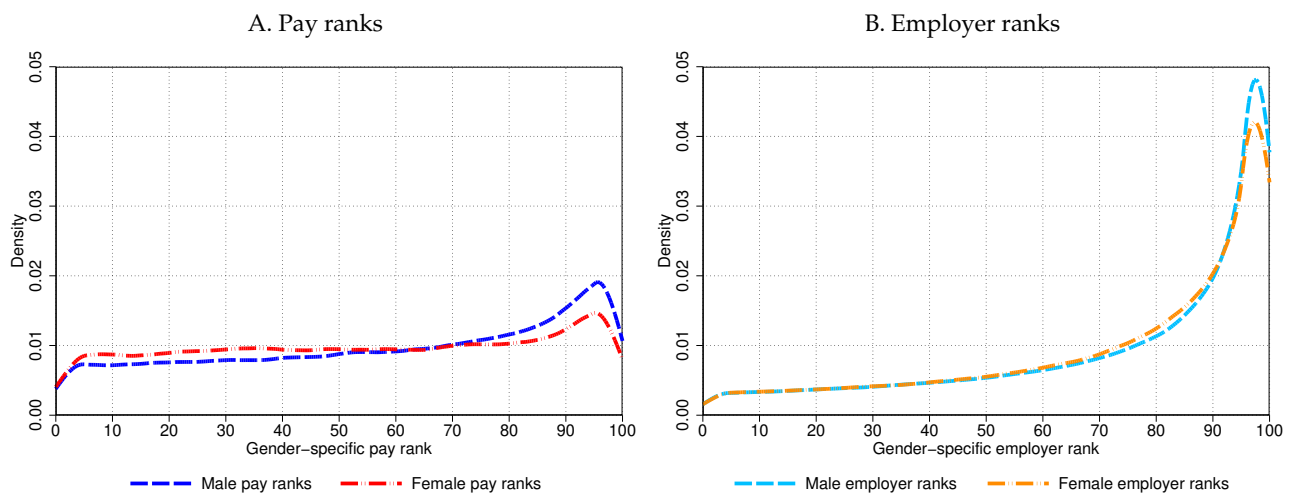
¹⁴Appendix B.9 shows that PageRanks strongly but imperfectly correlate with two other employer rank measures, namely the poaching rank ([Moscarini and Postel-Vinay, 2008](#); [Bagger and Lentz, 2019](#)) and the net poaching rank ([Haltiwanger et al., 2018](#); [Moscarini and Postel-Vinay, 2018](#)). An advantage of the PageRank is that it uses more information per worker transition in constructing an employer ranking, which reduces misclassifications of employer ranks.

¹⁵In all analyses, we study gender-specific employment-weighted estimates of the unweighted PageRanks as described above. PageRanks are not by construction restricted to having any particular employment-weighted mean value (e.g., 50) since the PageRank estimation is independent of the cross-sectional employment distribution.

Fact 1. *Employment is concentrated in high-ranked but not necessarily high-paying employers for both genders.*

Figure 2 compares the employment distributions of men and women across pay ranks and across employer ranks. Panel A shows employment is weakly positively related to pay for both genders. In contrast, Panel B shows that employment is strongly related to employer ranks for both genders. Furthermore, the employment distribution of women looks more similar over employer ranks than over pay ranks. Women’s mean employer pay rank is 53.9, whereas men’s is 58.3, implying a gender gap in pay ranks of 4.4 percentiles. On the other hand, women’s mean employer rank is 73.7, whereas men’s is 74.4, implying a gender gap in employer ranks of 0.7 percentiles. This fact suggests that nonpay job attributes are important for understanding labor market sorting of both men and women.

Figure 2. Densities over pay ranks and employer ranks, by gender




Note: Figure shows kernel density estimates of male and female employment distributions across establishments ranked by gender-specific pay FE estimates based on earnings equation (1) (Panel A) and ranked by gender-specific employer ranks using PageRank estimates from equation (4) (Panel B).

Fact 2. *Mean employer ranks are more increasing in pay ranks for men than for women.*

Figure 3 shows that employer ranks are positively correlated with pay ranks for both men and women. However, the employer rank-pay gradient is steeper for men, especially in the bottom half of the employer pay distribution. This result implies that there exist low-paying jobs that are relatively attractive for women but less so for men. Therefore, pay is a relatively more important determinant of the overall evaluation of an employer for men than for women.¹⁶ This fact suggests that there is differential sorting across genders by pay versus nonpay job attributes.

¹⁶Appendix B.10 presents several robustness checks. Figure B.21 shows the relationship between employer ranks and pay ranks across sectors. Table B.4 shows that this fact is not driven by sectoral or geographic differences. Table B.5 shows that this fact is consistent with the dynamics of pay for different worker transitions across employer ranks.

Figure 3. Employer ranks versus pay ranks, by gender



Note: Figure shows mean gender-specific employer ranks using PageRank estimates based on equation (4) across quantile bins of the distribution of gender-specific pay FE estimates based on earnings equation (1).

Fact 3. *Men and women agree more on pay and revealed-preference rankings of higher-ranked employers.*


Figure 4 shows that while pay and employer ranks are positively correlated across genders, there is substantial disagreement between men and women on employer pay ranks (Panel A) and revealed-preference ranks (Panel B). For example, the interquartile range of women's pay around the median pay for men spans the 45th to 65th percentiles of pay for women. Analogously, the interquartile range of women's employer ranks around the median rank for men spans the 48th to 78th percentiles of employer ranks for women. Furthermore, men and women agree notably more on the pay (ranks) of higher-paying (higher-ranked) employers.¹⁷ This fact suggests that barriers keeping women out of jobs that are also the most desirable to men are especially costly.

4 Model

The previous empirical facts motivate a structural model with the following ingredients. First and foremost, the model must allow an employer's revealed-preference rank to differ from its pay rank. To rationalize this, workers value an employer's *amenities* in addition to pay. Second, the model must generate differences in pay and amenities across employers. To rationalize this, the labor market is frictional. Third, the model must admit gender differences in pay, revealed-preference ranks, and employment within a firm. To rationalize this, employers post gender-specific wages, amenities, and job vacancies. We combine these elements in an equilibrium model of the labor market.

¹⁷For robustness, Appendix B.11 shows the same relation between female and male pay ranks (Figure B.22) and between female and male employer ranks (Figure B.23) by sector. The same qualitative conclusions apply within sectors.

Figure 4. Percentiles of female versus male employer pay and ranks



Note: Figure shows various percentiles of the employer pay FE distribution based on estimating earnings equation (1) for women against that for men (Panel A) and of the employer rank distribution using estimated PageR-ranks based on equation (4) for women against that for men (Panel B).

4.1 General Environment

A measure 1 of workers and measure E of firms meet in a continuous-time frictional labor market.

4.2 Workers

Workers are infinitely lived, risk neutral, and discount the future at rate ρ . A worker type is a duplet of ability $a \in [\underline{a}, \bar{a}]$ and gender $g \in \{M, F\}$ with measure $\mu_{a,g}$ such that $\sum_{g=M,F} \int_a \mu_{a,g} da = 1$. At any point in time, workers find themselves either employed or nonemployed.¹⁸

Job search. While nonemployed, workers receive flow utility $b_{a,g}$ and engage in random job search within segmented labor markets by worker type. Search is random in the sense that workers cannot direct their search to specific firms. Labor markets are segmented in the sense that workers search for jobs in a market specific to their type. While employed, workers receive flow utility $x = w + \pi$ equal to the sum of their wage, w , and job amenity value, π . Employed workers also engage in on-the-job search within the same segmented markets.

Workers receive regular job offers with arrival rate $\lambda_{a,g}^U$ from nonemployment and with rate $\lambda_{a,g}^E$ from employment. While regular job offers admit free disposal, workers also receive mandatory job

¹⁸Throughout the paper, we think of the “nonemployed” workers in the model as capturing the pool of workers who are not formally employed—that is, those who are unemployed, on temporary parental or other leave, marginally attached to the labor force, or in informal employment. In mapping the model to the data, our estimation of labor market parameters will take into account that some workers might spend longer periods outside of formal employment because of these factors.

offers (sometimes termed a “Godfather shock,” or an offer one cannot refuse) at rate $\lambda_{a,g}^G$ in both employment states. We think of the latter as capturing, among other things, spousal relocation problems and other idiosyncratic reasons for switching jobs. We write $\lambda_{a,g}^E = s_{a,g}^E \lambda_{a,g}^U$ and $\lambda_{a,g}^G = s_{a,g}^G \lambda_{a,g}^U$, where $s_{a,g}^E$ and $s_{a,g}^G$ are the relative hazards of regular and mandatory on-the-job offers, respectively.

A job offer is an opportunity to work at some firm with wage w and amenity value π drawn from a distribution $\tilde{F}(w, \pi)$, which workers take as exogenous but that is determined endogenously through firms’ equilibrium decisions. Since a worker’s state is summarized by their flow utility $x = w + \pi$, jobs are ranked on a ladder according to x , and we can restrict attention to the flow-utility offer distribution $F(x)$. Jobs are endogenously terminated when a worker with flow utility x accepts a higher-utility job at rate $\lambda_{a,g}^E(1 - F(x))$ and exogenously terminated when the worker relocates to a randomly drawn job at rate $\lambda_{a,g}^G$, or when moving into nonemployment at rate $\delta_{a,g}$. When setting wages, employers take into account that flow rates differ by gender—a phenomenon we term *statistical discrimination* (Arrow, 1971).

Value functions. The value of an employed worker of type (a, g) in a job with flow utility x is

$$\begin{aligned} \rho S_{a,g}(x) = & x + \lambda_{a,g}^E \int_{x' \geq x} [S_{a,g}(x') - S_{a,g}(x)] dF_{a,g}(x') + \lambda_{a,g}^G \int_{x'} [S_{a,g}(x') - S_{a,g}(x)] dF_{a,g}(x') \\ & + \delta_{a,g} [W_{a,g} - S_{a,g}(x)] . \end{aligned} \quad (5)$$

Analogously, the value of a nonemployed worker of type (a, g) is summarized as follows:

$$\rho W_{a,g} = b_{a,g} + (\lambda_{a,g}^U + \lambda_{a,g}^G) \int_{x'} \max \{S_{a,g}(x') - W_{a,g}, 0\} dF_{a,g}(x') . \quad (6)$$

Policy function. Strict monotonicity of $S_{a,g}(x)$ implies that optimal job acceptance of the nonemployed follows a threshold rule with reservation flow utility $\phi_{a,g}$. A nonemployed worker accepts an offer if $x \geq \phi_{a,g}$ and rejects it otherwise. The reservation flow utility simply equals the sum of the flow value of nonemployment plus the forgone option value of receiving job offers while nonemployed:

$$\phi_{a,g} = b_{a,g} + (\lambda_{a,g}^U - \lambda_{a,g}^E) \int_{x' \geq \phi_{a,g}} \frac{1 - F_{a,g}(x')}{\rho + \delta_{a,g} + \lambda_{a,g}^G + \lambda_{a,g}^E [1 - F_{a,g}(x')]} dx' . \quad (7)$$

Employed workers in a job with flow utility x simply accept any job that delivers flow utility $x' > x$.

Nonemployment and utility dispersion. Since in equilibrium no firm posts a contract worth less than $\phi_{a,g}$ in any market (a, g) , the steady-state nonemployment rate for each worker type is

$$u_{a,g} = \frac{\delta_{a,g}}{\delta_{a,g} + \lambda_{a,g}^U + \lambda_{a,g}^G}. \quad (8)$$

The cross-sectional distribution of flow utilities is given by

$$G_{a,g}(x) = \frac{F_{a,g}(x)}{1 + \kappa_{a,g}^E [1 - F_{a,g}(x)]},$$

where $\kappa_{a,g}^E = \lambda_{a,g}^E / (\delta_{a,g} + \lambda_{a,g}^G)$ governs the effective speed of workers climbing the job ladder.

4.3 Firms

Firms differ in four dimensions: productivity $p \in [p, \bar{p}] \subset \mathbb{R}_{++}$ as in [Burdett and Mortensen \(1998\)](#), a set of *gender wedges* $z_{a,g} \in [z, \bar{z}] \subset \mathbb{R}$ representing the firm's disutility from worker type (a, g) as in [Becker \(1971\)](#), a set of amenity cost shifters $c_{a,g}^{\pi,0} > 0$ as in [Hwang et al. \(1998\)](#), and a set of vacancy cost shifters $c_{a,g}^{v,0} > 0$. Thus, a firm's type is $j = (p, \{z_{a,g}\}_{a,g}, \{c_{a,g}^{v,0}\}_{a,g}, \{c_{a,g}^{\pi,0}\}_{a,g})$, which we assume is continuously distributed according to $\Gamma(j)$.

Wages, amenities, and job vacancies. Firms deliver value to workers through two channels. First, they post in each market a wage rate $w_{a,g}$. Second, they also post a market-specific value of amenities $\pi_{a,g}$.¹⁹ Following [Hwang et al. \(1998\)](#), we assume that the cost of producing amenities $\pi_{a,g}$ is paid per worker of type (a, g) employed at the firm and that per-worker amenity flow costs can be written as $c_{a,g}^\pi(\pi_{a,g}) = c_{a,g}^{\pi,0} \times \tilde{c}_{a,g}^\pi(\pi_{a,g})$, where $\tilde{c}_{a,g}^\pi(\cdot)$ satisfies $\tilde{c}_{a,g}^\pi(0) = 0$, $\partial \tilde{c}_{a,g}^\pi(0) / \partial \pi = 0$, and $\partial \tilde{c}_{a,g}^\pi(\pi) / \partial \pi, \partial^2 \tilde{c}_{a,g}^\pi(\pi) / \partial \pi^2 > 0$ for all $\pi > 0$ and all (a, g) . In order to recruit workers and produce output, firms post $v_{a,g}$ job vacancies in each market subject to flow cost $c_{a,g}^v(v_{a,g}) = c_{a,g}^{v,0} \times \tilde{c}^v(v_{a,g})$, where $\tilde{c}^v(\cdot)$ satisfies $\tilde{c}^v(0) = 0$, $\partial \tilde{c}^v(0) / \partial v = 0$, and $\partial \tilde{c}^v(v) / \partial v, \partial^2 \tilde{c}^v(v) / \partial v^2 > 0$ for all $v > 0$.

Production. A firm with productivity p employing $\{l_{a,g}\}_{a,g}$ workers of each type produces output according to the following linear production technology:

$$y(p, \{l_{a,g}\}_{a,g}) = p \sum_{g=M,F} \int_a a l_{a,g} da.$$

¹⁹Appendix C.4 presents an alternative model, in which firms produce an amenity vector with gender-specific utility weights, and establishes conditions for *observational equivalence* and *counterfactual equivalence* between the two models.

Gender wedge. In addition to output specified above, the model allows employers to care about employing different worker types. We model this as a set of gender wedges $\{z_{a,g}\}_{a,g}$, which, as two special cases, may capture taste-based discrimination as in [Becker \(1971\)](#) or firm-level comparative advantages in productivity across genders related to “brain versus brawn” ([Goldin, 1992](#); [Rendall, 2018](#)). We restrict gender wedges to the form $z_{a,g} = \mathbf{1}[g = F]z_a$, where z_a guides an employer’s relative preference for employing men over women among workers of ability a .

Value function. Firms post wages, amenities, and vacancies in each market to maximize steady-state flow payoff. The value $\Pi(j)$ of a firm of type $j = (p, \{z_a\}_a, \{c_{a,g}^{v,0}\}_{a,g}, \{c_{a,g}^{\pi,0}\}_{a,g})$ is given by

$$\rho\Pi(j) = \max_{\{w_{a,g}, \pi_{a,g}, v_{a,g}\}_{a,g}} \left\{ \sum_{g=M,F} \int_a \left[pa - w_{a,g} - c_{a,g}^{\pi}(\pi_{a,g}) - z_{a,g} \right] l_{a,g}(w_{a,g}, \pi_{a,g}, v_{a,g}) - c_{a,g}^v(v_{a,g}) da \right\}. \quad (9)$$

4.4 Matching

The effective mass of job searchers in market (a, g) equals

$$U_{a,g} = \mu_{a,g} \left[u_{a,g} + s_{a,g}^E(1 - u_{a,g}) + s_{a,g}^G \right], \quad \forall(a, g). \quad (10)$$

The total mass of vacancies posted in market (a, g) across firm types j equals

$$V_{a,g} = E \int_j v_{a,g}(j) d\Gamma(j), \quad \forall(a, g). \quad (11)$$

A Cobb-Douglas matching function with constant returns to scale combines the mass of job searchers with the mass of job vacancies to produce matches between workers and firms, $m_{a,g}$, according to

$$m_{a,g} = \chi_{a,g} V_{a,g}^\alpha U_{a,g}^{1-\alpha}, \quad \forall(a, g),$$

where $\chi_{a,g} > 0$ is the matching efficiency and $\alpha \in (0, 1)$ is the matching elasticity with respect to aggregate vacancies. Define labor market tightness as

$$\theta_{a,g} = \frac{V_{a,g}}{U_{a,g}}, \quad \forall(a, g). \quad (12)$$

The job-finding rates among nonemployed workers, $\lambda_{a,g}^U$; the job-finding rates among the employed, $\lambda_{a,g}^E$; the arrival rates of mandatory offers, $\lambda_{a,g}^G$; and firms' job-filling rates, $q_{a,g}$, are given by

$$\lambda_{a,g}^U = \chi_{a,g} \theta_{a,g}^\alpha, \quad \lambda_{a,g}^E = s_{a,g} \lambda_{a,g}^U, \quad \lambda_{a,g}^G = s_{a,g}^G \lambda_{a,g}^U \text{ and } q_{a,g} = \chi_{a,g} \theta_{a,g}^{\alpha-1}, \quad \forall (a, g). \quad (13)$$

4.5 Firm Size Distribution

The following Kolmogorov forward (or Fokker-Planck) equation describes the law of motion of firm sizes given a firm's flow-utility and vacancy policy (x, v) , the market distribution of flow utilities $F_{a,g}(x)$, and market tightness $\theta_{a,g}$:

$$\dot{l}_{a,g}(x, v) = \left[-\delta_{a,g} - \lambda_{a,g}^E [1 - F_{a,g}(x)] - \lambda_{a,g}^G \right] l_{a,g}(x, v) + \left[\frac{u_{a,g} + (1 - u_{a,g}) s_{a,g}^E G_{a,g}(x) + s_{a,g}^G}{u_{a,g} + (1 - u_{a,g}) s_{a,g}^E + s_{a,g}^G} \right] v q_{a,g}.$$

Solving for the stationary firm size distribution, we find

$$l_{a,g}(x, v) = \left(\frac{1}{\delta_{a,g} + \lambda_{a,g}^G + \lambda_{a,g}^E [1 - F_{a,g}(x)]} \right)^2 \frac{v}{V_{a,g}} \mu_{a,g}(u_{a,g} + s_{a,g}^G) \lambda_{a,g}^U (\delta_{a,g} + \lambda_{a,g}^G + \lambda_{a,g}^E). \quad (14)$$

4.6 Equilibrium Characterization

in Appendix C.1, we define a *stationary equilibrium* of the economy. The assumed market segmentation and linearity of the production technology allow us to keep this problem tractable in spite of the many dimensions of worker and firm heterogeneity. These assumptions allow us to divide the firm's problem into separate subproblems by market. Conditional on productivity, a firm's optimal choice in each market is essentially independent of all other markets, which means that we can solve the firm's problem in each market in isolation.

For any posted wage-amenity combination, firms find themselves ranked on a market-specific ladder according to their flow-utility offer x . An argument analogous to that in [Burdett and Mortensen \(1998\)](#) shows that the equilibrium offer distribution $F_{a,g}(x)$ and the cross-sectional distribution $G_{a,g}(x)$ are continuous and strictly increasing for $x > \max\{\underline{p}a - \mathbf{1}[g = F]\bar{z}, \phi_{a,g}\}$ up to some maximum value in each market (a, g) . Next, we characterize firms' optimal policy functions.

Lemma 1 (Optimal amenities). *A firm's optimal amenity policy $\pi_{a,g}^*(\cdot)$ is strictly decreasing in its amenity cost shifter $c_{a,g}^{\pi,0}$ and is invariant to all other parameters. Furthermore, $0 < c_{a,g}^{\pi}(\pi_{a,g}^*) < \pi_{a,g}^*$.*

Proof. See Appendix C.2.1. □

Lemma 1 extends to our setting a key result in Hwang et al. (1998), who also assume that firms are heterogeneous in their convex-increasing per-worker cost of amenities. Intuitively, firms optimally offer amenities up to the point where the marginal cost of amenities equals that of wages, which equals one. That the cost-minimization problem does not depend on a firm's productivity, gender wedge, or recruiting costs follows from two assumptions: that worker utility is additively separable between wages and amenities and that the amenity cost is paid per worker. An implication of Lemma 1 is that because of the bijection between firm-specific amenity cost shifters and optimal amenity values, we can treat $\pi_{a,g}^*$ as an exogenous firm-level parameter. Furthermore, in model counterfactuals, a firm's optimal amenity choice remains at the estimated value unless there are changes to its amenity cost function relative to its wage cost function.

Define a firm's *composite productivity* in market (a, g) as $\tilde{p}_{a,g} = pa + \pi_{a,g} - c_{a,g}^\pi(\pi_{a,g}) - z_{a,g}$. We can treat $\tilde{p}_{a,g}$ as an exogenous firm characteristic, allowing us to rewrite the problem of a firm as

$$\rho\Pi_{a,g}(\tilde{p}_{a,g}, c_{a,g}^{v,0}) = \max_{x,v} \left\{ [\tilde{p}_{a,g} - x] l_{a,g}(x, v) - c_{a,g}^v(v) \right\}, \quad \forall(a, g). \quad (15)$$

Therefore, the current model is essentially isomorphic to one without amenities or gender wedges, but with two modifications.²⁰ First, productivity p is replaced by composite productivity \tilde{p} . Second, wages w are replaced by flow utility x . This isomorphism allows us to derive comparative statics with respect to the different components of $\tilde{p}_{a,g}$.

Lemma 2 (Optimal market selection). *A firm optimally employs workers in market (a, g) if $\tilde{p}_{a,g} > \phi_{a,g}$.*

Proof. See Appendix C.2.2. □

A firm makes positive monetary profits if $pa + \pi_{a,g} - c_{a,g}^\pi(\pi_{a,g}) > \phi_{a,g}$.²¹ However, Lemma 2 states that, due to the presence of gender wedges, this condition is neither necessary nor sufficient for a firm to select into a market. Depending on z_a in relation to the monetary surplus $pa + \pi_{a,g} - c_{a,g}^\pi(\pi_{a,g}) - \phi_{a,g}$ in each market, the firm may hire any combination of genders: both, either one, or none (in which case it does not operate).

Lemma 3 (Optimal vacancies). *A firm's optimal vacancy policy $v_{a,g}^*(\cdot)$ is strictly increasing in productivity p , strictly decreasing in the vacancy cost shifter $c_{a,g}^{v,0}$ for all worker types, and strictly decreasing (constant) in z_a for women (men).*

²⁰See Engbom and Moser (2018) for an example of such a model.

²¹We implicitly assume here that a firm's productivity is high enough to pay the minimum wage in monetary units.

Proof. See Appendix C.2.3. □

The intuition behind Lemma 3 is that higher-productivity firms have a higher marginal payoff per contacted worker; thus, they invest more into recruiting both men and women. The opposite is true for female vacancies at firms with a higher gender wedge in their payoff function. Naturally, firms with a higher vacancy cost post fewer vacancies for both genders.

Lemma 4 (Optimal flow utility and wages). *A firm's optimal flow-utility policy $x_{a,g}^*(\cdot)$ and wage policy $w_{a,g}^*(\cdot)$ are strictly increasing in p for all worker types, constant in the vacancy cost shifter $c_{a,g}^{v,0}$ for all worker types, and strictly decreasing (constant) in the gender wedge z_a for women (men).*

Proof. See Appendix C.2.4. □

Lemma 4 extends the comparative statics results with respect to wages in Mortensen (2003) to an environment with richer employment contracts (amenities and wages, instead of just wages) and richer sources of worker mobility (Godfather shocks and heterogeneous arrival rates from nonemployment and employment, instead of just homogeneous arrival rates). Intuitively, firms with a larger payoff from employing a given worker optimally offer workers higher utility through wages in order to attract and retain a larger workforce.

Lemma 5 (Optimal employment). *A firm's optimal employment $l_{a,g}^*(\cdot)$ is strictly increasing in p for all worker types, strictly decreasing in the vacancy cost shifter $c_{a,g}^{v,0}$ for all worker types, and strictly decreasing (constant) in the gender wedge z_a for women (men).*

Proof. See Appendix C.2.5. □

Lemma 5 states that firms with higher composite productivity \tilde{p} have greater steady-state employment, which is a combination of their rank in the job ladder, as guided by their flow-utility rank, and their recruitment intensity, as guided by the share of their aggregate share of vacancies.

4.7 Equilibrium Wage Equation

Our equilibrium model provides a microfoundation for the decomposition of log wages into worker FEs and gender-specific employer FEs from Section 3.1, building on Card et al.'s (2016) variant of the framework developed by AKM. To demonstrate this, we provide a set of sufficient conditions for the log-wage decomposition to obtain as an equilibrium outcome in the model.

Assumption 1 (Vacancy cost function). *Vacancy-posting costs $c_{a,g}^{v,0}$ scale linearly in worker ability a :*

$$c_{a,g}^{v,0} = ac_g^{v,0}, \quad \forall a.$$

Assumption 1 could reflect the payment of recruiting costs in terms of time given to new hires for orientation and training, or in terms of the time of equally skilled workers devoted to recruiting.

Assumption 2 (Job offer arrival and separation rates). *The job offer arrival rate $\lambda_{a,g}^U$, relative arrival rates of optional job offers $s_{a,g}^E$, relative arrival rates of mandatory job offers $s_{a,g}^G$, and separation rates $\delta_{a,g}$ are constant in worker ability a :*

$$\lambda_{a,g}^U = \lambda_g^U, \quad s_{a,g}^E = s_g^E, \quad s_{a,g}^G = s_g^G, \quad \delta_{a,g} = \delta_g, \quad \forall a.$$

Assumption 2 allows for differential worker mobility across, but not within, genders.

Assumption 3 (Amenity cost function). *The amenity creation cost function $c_{a,g}^\pi(\pi)$ takes on the following piece-rate form:*

$$c_{a,g}^\pi(\pi) = ac_g^{\pi,0} \tilde{c}\left(\frac{\pi}{a}\right), \quad \forall a.$$

Assumption 3 states that the cost of creating amenities is proportional to worker ability and that amenities are paid to a worker as a piece rate in his or her ability. A natural interpretation for this would be that some amenities involve time spent off work, such as paid parental leave. In this case, the cost of providing some units of time to a worker in amenities scales linearly in the worker's ability or forgone production due to the worker's absence from the job.

Assumption 4 (Flow values of nonemployment and gender wedges). *The flow values of nonemployment $b_{a,g}$ and gender wedges z_a scale linearly in worker ability a :*

$$b_{a,g} = b_g a, \quad z_a = z a, \quad \forall a.$$

Assumption 4 ensures symmetric participation and composite productivity across markets. It can be justified by higher-ability workers' also having higher returns to nonemployment—for example, because they are more skilled at home production or in informal employment—and by employers willingness to give up a fraction of workers' output to avoid interacting with them.

The following result links the structural model to the reduced-form approach in Section 3.1.

Proposition 1 (Equilibrium Wage Equation). *Under Assumptions 1–4, the equilibrium wage of a worker with ability a and gender g at a firm with composite productivity \tilde{p}_g and amenity cost shifter $c_g^{\pi,0}$ is*

$$\ln w_{a,g}(\tilde{p}_g, c_g^{\pi,0}) = \underbrace{\alpha_a}_{\text{“worker FE”}} + \underbrace{\psi_g(\tilde{p}_g, c_g^{\pi,0})}_{\text{“gender-firm FE”}}, \quad (16)$$

where

$$\alpha_a = \ln a, \quad \psi_g(\tilde{p}_g, c_g^{\pi,0}) = \ln \left(\tilde{p}_g - \pi_g^*(c_g^{\pi,0}) - \int_{\tilde{p}' \geq \phi_g}^{\tilde{p}_g} \left[\frac{1 + \kappa_g^E [1 - F_g(x_g^*(\tilde{p}_g))]}{1 + \kappa_g^E [1 - F_g(x_g^*(\tilde{p}'))]} \right]^2 d\tilde{p}' \right). \quad (17)$$

Proof. See Appendix C.2.6. □

Proposition 1 shows that under appropriate scaling assumptions, equilibrium wages in the model are log-additive between a worker component (“worker FE”) and a gender-specific firm component (“gender-specific firm FE”). The worker FE α_a is a strictly monotonic transformation of worker ability. The gender-specific firm FE $\psi_g(\tilde{p}_g, c_g^{\pi,0})$ depends only on parameters that are gender-firm specific—namely, a firm’s composite productivity \tilde{p}_g and its amenity cost shifter $c_g^{\pi,0}$. Therefore, the equilibrium model provides a microfoundation for the wage equation with gender-specific employer pay components used in our empirical investigation. For the remainder of the analysis, we maintain assumptions 1–4 and focus on differences in gender-specific firm FEs between men and women.²²

4.8 Equilibrium Properties

The above model has three notable equilibrium properties. First, the model can rationalize job-to-job transitions with wage declines. On one hand, workers receive exogenous relocation shocks that result in forced transitions from wage w to $w' < w$. On the other hand, workers may endogenously transition from wage-utility combination (w, x) to (w', x') with $x' > x$ but $w' < w$.

Second, “discriminatory” firms (as captured by the gender wedge z) can survive in a frictional environment.²³ A prediction of Becker’s (1971) seminal framework of taste-based discrimination is

²²Under these assumptions, Appendix C.3 demonstrates log-additivity in worker ability of key model objects: workers’ value functions, workers’ policy function, and firms’ value function. We thank Thibaut Lamadon for suggesting this to us.

²³We do not want to claim that the gender wedge z relates only to discrimination. On the contrary, we think of it as capturing many different mechanisms. Among such mechanisms, taste-based discrimination is of particular interest. All else being equal, higher taste-based discrimination against female workers is associated with higher values of z .

that in a competitive market, employers with a distaste for certain workers are driven out of the market. In contrast, in the current model, firms with a nonzero gender wedge z can survive in the presence of labor market frictions.

Third, even “nondiscriminatory” firms (or those with $z = 0$) may pay women less, owing to statistical discrimination based on gender-specific transition rates, compensating differentials, or their equilibrium response to the presence of other discriminatory employers.

4.9 Discussion of Model Assumptions

We now turn to a brief discussion of some of the more restrictive modeling assumptions and their implications. We first assume that output is linear within and additively separable across worker types. These assumptions allow for considerable analytical tractability, but we argue that they are also in line with our empirical evidence.

That output is linear within worker types is not particularly restrictive, since a firm’s net payoff function is already concave because of the convex vacancy cost. Conceptually, there is no reason not to simultaneously allow for curvature in the ability-weighted number of workers of each type. However, if the marginal product of a given worker type were exceedingly high for small numbers of workers—as would be the case with standard constant-elasticity-of-substitution specifications—then we would see every firm employing a strictly positive mass of each worker type. This outcome would be clearly at odds with the presence of single-gender firms in the data.

Additive separability of output in worker types allows the model to admit a log-linear wage equation, which we require to take the model to the data. Assuming complementarities between genders would lead to the counterfactual implication that no single-gender firms could exist. Supporting our assumption, [Fukui et al. \(2020\)](#) find a small crowd-out between women and men in the U.S. Therefore, a natural starting point treats men and women as perfect substitutes in production.

A second assumption revolves around labor market segmentation by worker types, which allows firms to tailor wages, amenities, and vacancies to each market. This assumption considerably simplifies the analysis of the firm’s problem. We argue that the assumption is both reasonable theoretically and also in line with salient empirical patterns.

That firms can direct wages and vacancies toward certain worker types may seem at odds with nondiscrimination laws. But of course a firm need not publicly post different wages or job openings by gender in order to discriminate. Such differences may arise in more subtle ways during résumé screenings, during interviews, and at the negotiation table ([Goldin and Rouse, 2000](#)). This assumption

is consistent with evidence that men and women differentially apply to jobs based on the wording of vacancies (Abraham and Stein, 2020) and accept jobs subject to deadlines (Cortés et al., 2020).

Empirically, there are also good reasons to adopt market segmentation. First and foremost, our model must confront the significant differences in the gender-specific employer component of pay, amenity utilization, and employment of men and women within the same employer in the Brazilian data. In the previous section, we have already documented gender differences in pay and employment. In Appendix E.4, we show that there are also large differences in amenity utilization across genders within employers. Our model provides a natural way of rationalizing these differences.

For robustness, we relax some of the above assumptions by solving three alternative models. We find these alternative models to have salient counterfactual predictions. A model with firms offering a single wage for men and women fails to account for the empirical within-firm pay differences documented in section 3.1. A model in which amenity values are the same across genders within a firm counterfactually predicts no dispersion in firm ranks conditional on gender-specific pay, while one in which firms produce an amenity vector with gender-specific utility weights is discussed in Appendix C.4. Finally, a model in which vacancy costs are over the sum of gender-specific vacancies or in which vacancies are not directed toward genders fails to account for the empirical dispersion of female employment shares and hence the between-employer pay gap in the data.²⁴

Furthermore, we treat workers with and without children the same, regardless of their parental status. Importantly, though, our model allows for gender-specific differences in worker flow rates. This is motivated by the empirical evidence in Appendix B.8 as well as by the argument that even nonparent women may be treated differently by employers, in anticipation of future childbirths.²⁵

Finally, the assumptions underlying Proposition 1 allow us to make significant progress in bringing the model to the data. They provide a microfoundation for the piece-rate wage equation that is commonly assumed in other models (e.g., Sorkin, 2018). If these assumptions do not hold exactly in the data—for example when labor market parameters differ across ability levels—then the decomposition in equation (16) will not hold exactly. However, this turns out to be less problematic quantitatively, as Engbom and Moser (2018) show that an AKM decomposition of log wage predicts around 99 percent of the variance of log earnings in data simulated from a model that allows for flexible

²⁴We later solve the model with gender-neutral wage offers when simulating the effects of an equal-pay policy. The other two models are solved in Appendix C.5. The model with directed vacancy posting and a joint cost function in Appendix C.5.1 predicts that with the exception of knife-edge cases, there exist no dual-gender firms. The model with undirected vacancy posting in Appendix C.5.2 predicts that, quantitatively, there is far too little dispersion in female employment shares compared with the empirical distribution we see in the data.

²⁵It would be straight forward to re-estimate our model separately for ever-parent and never-parent workers and by different education groups to account for within-gender differences in sorting.

variation in labor market parameters across ability types estimated to the same data from Brazil.

5 Identification Strategy

To bridge the model and the data, we connect key model objects with their empirical counterparts. Our starting point is the special case of the model characterized in Proposition 1 of Section 4.7. Under the maintained assumptions, this case allows us to pool workers of different ability types in the data and drop a from all this section’s subscripts. We adopt a three-step identification strategy.

5.1 Step 1: Employer Ranks

In the first step, we estimate revealed-preference ranks of employers by gender, using the PageRank index (Page et al., 1998; Sorkin, 2018) described in Section 3.2. This constitutes a set of $N_M + N_F$ estimates, where N_M and N_F are the empirical numbers of establishments hiring men and women, respectively. The PageRank index represents the asymptotic share of time a representative worker would spend at a given employer. This notion of employer rank coincides with that in the structural model of Section 4, in which workers are less likely to endogenously separate from and more likely to accept offers at higher-utility employers. In what follows, we conflate ranks and employer identities by indexing establishments by their rank $r_g \in \{1, 2, \dots, R_g\}$, where 1 is the lowest and R_g is the highest rank for workers of gender g .

5.2 Step 2: Labor Market Parameters

In the second step, we estimate labor market parameters by combining employer ranks from Step 1 with monthly information on worker flows.²⁶ We seek gender-specific estimates of the cumulative density function (CDF) of offers F_g^r , separation rates δ_g , job-finding rates from nonemployment λ_g^U , the relative arrival rate of mandatory on-the-job offers s_g^G , and relative arrival rates of voluntary on-the-job offers s_g^E .²⁷ This constitutes a set of $N_M + N_F + 8$ parameters. To this end, we exploit the model’s job-ladder property that worker transitions depend only on ordinal employer ranks.

²⁶The high-frequency nature of our data allows us obtain more precise estimates of employer ranks than has been possible in previous work. For example, Sorkin (2018) uses quarterly data to compute employer ranks based on what is effectively annual information on employment spells. We find that time aggregation bias (Moscarini and Postel-Vinay, 2018) can be substantial when using aggregated data at the quarterly or annual level to repeat our estimates.

²⁷Recall that we map the “nonemployed” workers in the model into the pool of workers who are not formally employed—that is, those who are unemployed, on temporary parental or other leave, marginally attached to the labor force, and in informal employment.

Job offer distributions. After ordering employers by their revealed-preference rank r , we compute the share of hires from nonemployment of each employer j out of total hires from nonemployment to estimate the gender-specific offer CDF $F_g^r = F_g(x_g^r)$.

Exogenous separation rates. We identify δ_g off separation rates into nonemployment:

$$\hat{\delta}_i = \mathbb{E}_i \mathbf{1} \left[\text{nonemployed}_{i,t+1} \mid \text{employed}_{i,t}, \text{gender}_i = g \right].$$

Offer rates from nonemployment. We identify λ_g^U off a log-hazard model for the time it takes for a worker to return to the data from nonemployment:

$$\hat{\lambda}_g^U = 1 - \exp \left(\frac{\ln (\mathbb{E}_i \mathbf{1} [\text{nonemployment duration}_i \geq t \mid \text{gender}_i = g])}{t} \right).$$

Mandatory on-the-job offer rates. Two insights allow us to use information on worker transitions between employers to identify λ_g^G . First, we focus on transitions in rank space, not pay space. Second, the share of rank-increasing transitions due to mandatory on-the-job offers declines in F_g^r . Formally, the total number of job-to-job transitions from employer rank r is

$$J2J_g^r = n_g^r [\lambda_g^E (1 - F_g^r) + \lambda_g^G], \quad (18)$$

where n_g^r is the number of workers of gender g at r . Rearranging and taking expectations, we have

$$\hat{\lambda}_g^G = \mathbb{E}_i \left[\frac{J2J_g^{r\downarrow}}{n_g^r \hat{F}_g^r} \right],$$

where $J2J_g^{r\downarrow} = J2J_g^r - n_g^r (\lambda_g^E + \lambda_g^G) (1 - F_g^r)$ is the number of job-to-job transitions to lower ranks.

Voluntary on-the-job offer rates. On-the-job offers not associated with mandatory transitions must have been voluntary. Hence, once we know λ_g^G , we can use equation (18) to estimate λ_g^E as

$$\hat{\lambda}_g^E = \frac{J2J_g^r / n_g^r - \hat{\lambda}_g^G}{1 - \hat{F}_g^r}.$$

5.3 Step 3: Employer-Level Parameters and Values of Nonemployment

In the third step, we estimate employer-level parameters—productivity, amenity cost shifters, gender wedges, and vacancy cost shifters—together with workers’ flow values of nonemployment, using information on gender-specific employer ranks, pay, and labor market parameters. This constitutes a set of $3(N_M + N_F) + 2$ parameters. The fact that amenity values act as a residual between employers’ rank and pay allows us to set-identify amenity values for each employer.²⁸ We use structural equilibrium restrictions to further narrow the identified sets. Within this set, we pick an amenity vector that minimizes utility dispersion. Appendix D.1 presents an illustrative example of the identification routine with three employers. We now delineate the general case.

Using Lemma 1, we can search for amenity values rather than amenity cost shifters, since the two are isomorphic. With a slight abuse of notation, we will conflate (estimates of) the gender-specific employer component of pay in the data, denoted by ψ_j^g in the empirical earnings equation (1), with the theoretical notion of a gender-specific employer component of pay, denoted by $\psi_g(\tilde{p}_g, c_g^{\pi,0})$ in Proposition 1 of the equilibrium model. For the remainder of this section, we will denote the exponential of both by ψ_g^r for gender g at employer rank r and sometimes loosely refer to them as wages. Given wages $(\psi_g^1, \psi_g^2, \dots, \psi_g^{R_g}) \in \mathbb{R}_{++}^{R_g}$, the problem is to find separately by gender g a vector of amenity values $(\pi_g^1, \pi_g^2, \dots, \pi_g^{R_g}) \in \mathbb{R}_+^{R_g}$ subject to a sequence of flow-utility monotonicity constraints from Lemma 4:

$$\psi_g^r + \pi_g^r \leq \psi_g^{r+1} + \pi_g^{r+1}, \quad \forall r < R_g. \quad (19)$$

We pick an amenity vector from the identified set by minimizing the sum of squared differences between rank-adjacent utilities defined as²⁹

$$\sum_r \left[\left(\psi_g^{r+1} + \pi_g^{r+1} \right) - \left(\psi_g^r + \pi_g^r \right) \right]^2. \quad (20)$$

In partial equilibrium, the amenity values that are jointly consistent with employer ranks and pay form an identified set. We can do better, however, by additionally imposing structural equilibrium

²⁸The reason for set identification (as opposed to point identification) is that it is impossible to deduce cardinal utility measures from just ordinal employer rank and pay information absent additional restrictions on the environment.

²⁹In Appendix D.3, we report results from Monte Carlo simulations in support of this estimation routine. We tried several alternative ways of choosing amenities from the identified set in Monte Carlo simulations and found that choosing the utility-distance-minimizing performed best across different parameterizations of the data generating process.

restrictions on the amenity vector from Lemma 4:

$$\tilde{p}_g^r \leq \tilde{p}_g^{r+1}, \quad \forall r < R_g. \quad (21)$$

Rewriting firms' first-order condition (FOC) with respect to flow utility x in equation (15) yields³⁰

$$\tilde{p}_g^r = \psi_g^r + \pi_g^r + \frac{1 + \kappa_g^E(1 - F_g(x_g^r))}{2\kappa_g^E f_g(x_g^r)}, \quad \forall g. \quad (22)$$

To summarize, given wages $(\psi_g^1, \psi_g^2, \dots, \psi_g^{R_g})$, estimates of the offer distribution \hat{F}_g^r , and estimates of labor market parameters $\hat{\kappa}_g^E$, we find gender-specific amenity values $(\pi_g^1, \pi_g^2, \dots, \pi_g^{R_g})$ that minimize equation (20) subject to the constraints in equations (19), (21), and (22):

$$\begin{aligned} (\hat{\pi}_g^1, \hat{\pi}_g^2, \dots, \hat{\pi}_g^{R_g}) &= \arg \min_{(\pi_g^1, \pi_g^2, \dots, \pi_g^{R_g}) \in \mathbb{R}_+^{R_g}} \sum_r \left[(\psi_g^{r+1} + \pi_g^{r+1}) - (\psi_g^r + \pi_g^r) \right]^2 \\ \text{s.t. } \psi_g^r + \pi_g^r &\leq \psi_g^{r+1} + \pi_g^{r+1}, \quad \forall r < R_g \\ \psi_g^r + \pi_g^r + \frac{1 + \hat{\kappa}_g^E(1 - \hat{F}_g^r)}{2\hat{\kappa}_g^E \hat{f}_g^r} &\leq \psi_g^{r+1} + \pi_g^{r+1} + \frac{1 + \hat{\kappa}_g^E(1 - \hat{F}_g^{r+1})}{2\hat{\kappa}_g^E \hat{f}_g^{r+1}}, \quad \forall r < R_g. \end{aligned} \quad (23)$$

Given amenity estimates, we back out amenity cost shifters $\{\hat{c}_g^{\pi,0,r}\}$ from the functional form of the amenity cost function $\tilde{c}(\cdot)$. We combine estimates of amenity values, wages, and labor market parameters to back out composite productivities using equation (22). The definition of composite productivity for men yields employer productivity $\hat{p}^r = \hat{p}_M^r - \hat{\pi}_M^r + c_M^{\pi,r}(\hat{\pi}_M^r)$. For dual-gender employers, we estimate gender wedges as $\hat{z}^r = \hat{p}^r - \hat{p}_F^r + \hat{\pi}_F^r - c_F^{\pi,r}(\hat{\pi}_F^r)$.³¹ By the definition of the offer distribution, $v_g^r = f_g^r V_g$. Rearranging the FOCs for optimal vacancies, we estimate vacancy cost shifters as

$$\hat{c}_g^{v,0,r} = \frac{T_g(\hat{p}_g^r - \hat{x}_g^r)}{\partial \tilde{c}^v(\hat{f}_g^r V_g) / \partial v_g^r \left(\hat{\delta}_g + \hat{\lambda}_g^G + \hat{\lambda}_g^E(1 - \hat{F}_g^r) \right)^2}, \quad \forall r, \quad (24)$$

where $\hat{x}_g^r = \hat{w}_g^r + \hat{\pi}_g^r$ and $T_g = \mu_g[(u_g + s_g^G)\lambda_g^U(\delta_g + \lambda_g^G + \lambda_g^E)]/V_g$. Equation (24) relates the vacancy cost shifter $\hat{c}_g^{v,0,r}$ to the aggregate mass of vacancies V_g . Finally, gender-specific outside option values are estimated as $\hat{\phi}_g = \min_r \{\psi_g^r + \hat{\pi}_g^r\}$. Together with a value of the exogenous discount rate ρ , equation (7) yields estimates of the gender-specific flow values of nonemployment, \hat{b}_g .

³⁰See Appendix D.2 for details on the change of variables from r to x .

³¹Alternative normalizations of gender-specific employer FEs, which we have experimented with following the discussion in Section 3.1, load primarily onto estimates of the gender wedges.

To summarize, we have estimated amenity and vacancy cost shifters $\{\hat{c}_g^{\pi,0,r}, \hat{c}_g^{\pi,0,r}\}_r$ for each gender at every employer, productivities and gender wedges $\{\hat{p}^r, \hat{z}^r\}_r$ for dual-gender firms, a set of productivities $\{\hat{p}^r\}_r$ for firms employing only men, a set of female composite productivities $\{\hat{p}_F^r\}_r$ for firms employing only women, and gender-specific flow values of nonemployment \hat{b}_g .

Two final comments are in order. First, although we estimate relative amenity values across employers, we are unable to identify the mean of amenities, and hence the level of utility, for either gender. The reason is simply the invariance of revealed preferences to a level shift in utilities. Thus, we normalize amenities to be weakly positive for both genders. Second, we obviously cannot identify parameters relating both genders within an employer that hires workers of only one gender. We still use single-gender employers in the estimation, since they add to the identification of all other parameters. We assume that the parameters of single-gender employers with no workers of gender g are such that their composite productivity falls short of the outside option value for that gender, $\tilde{p}_g < \phi_g$. We keep these employers unchanged in counterfactuals involving gender-specific parameters.³²

6 Estimation Results

6.1 Exogenous Parameters and Functional Form Assumptions

We assume that the cost functions for amenities and vacancies are of the power form $\tilde{c}_g^\pi(\pi) = \pi^{\eta_\pi} / \eta_\pi$ with $\eta_\pi = 2$ and $\tilde{c}_g^v(v) = v^{\eta_v} / \eta_v$ with $\eta_v = 2$. Neither of these assumptions is relevant for model fit, as we can match the distributions of π_g^r and f_g^r establishment by establishment in the data regardless of functional forms or parameter values.³³ Finally, we assume a discount factor of $\rho = 0.051$, which corresponds to an annual compound real interest rate of 5.3 percent. See Table 4.

Table 4. Exogenous parameters

Parameter	Description	Value
η_π	Amenity cost elasticity	2.000
η_v	Vacancy cost elasticity	2.000
ρ	Discount rate	0.051

Note: Table shows parameter values that are exogenously set and treated as fixed.

³²We have repeated all of our simulations for the subset of dual-gender employers with similar results.

³³We have experimented with different elasticities of the amenity cost and vacancy cost functions for counterfactuals.

6.2 Labor Market Parameters and Flow Values of Nonemployment

Estimated labor market parameters are shown in Table 5. Women exhibit lower transition rates in general, both between employment states and between jobs. The implied nonemployment rates are $u_M = 0.243$ and $u_F = 0.244$, reflecting the presence of a large informal sector for both men and women in Brazil. While women are more likely to be permanently employed in the informal sector, men and women are similarly attached to the formal sector conditional on ever participating. For both men and women, mandatory on-the-job offers are about twice as frequent as voluntary ones. This finding may seem surprising but can be rationalized by noting two points. First, mandatory and voluntary job offers are close substitutes further down the job ladder, where many workers find themselves. Second, our estimates pertain to employer rankings in utility space and not in wage space as in traditional estimates (e.g., Jolivet et al., 2006). We also find that the flow value of nonemployment is higher for men than for women. In part, these estimates reflect the widespread presence of unregistered employment in Brazil, with workers of both genders frequently switching between formal and informal jobs (Meghir et al., 2015).

Table 5. Job offer arrival rates, job destruction rates, and flow values of nonemployment

Parameter	Description	Value	Implied rate
λ_M^u	Offer arrival rate from nonemployment (M)	0.100	0.100
λ_F^u	Offer arrival rate from nonemployment (F)	0.087	0.087
δ_M	Job destruction rate (M)	0.036	0.036
δ_F	Job destruction rate (F)	0.031	0.031
s_M^e	Relative arrival rate of voluntary on-the-job offers (M)	0.057	0.006
s_F^e	Relative arrival rate of voluntary on-the-job offers (F)	0.061	0.005
s_M^G	Relative arrival rate of mandatory on-the-job offers (M)	0.119	0.012
s_F^G	Relative arrival rate of mandatory on-the-job offers (F)	0.107	0.009
b_M	Flow value of nonemployment (M)	1.357	
b_F	Flow value of nonemployment (F)	1.267	


Note: Table shows parameter values that are estimated based on monthly job flows data. “M” denotes parameter for men; “F” denotes parameter for women. All implied rates are monthly.

6.3 Distributions of Productivity, Gender-Specific Amenities, and Wedges

Figure 5 shows the marginal distributions of estimated productivity, amenity values, and gender wedges. Productivity dispersion is substantially larger than that of amenities or gender wedges. Productivity and gender wedges are positively skewed. Male and female amenities are similarly dispersed and left skewed and have thinner tails than a normal distribution. An advantage of our es-

timination approach is that we do not impose ex ante restrictions on the distributions of gender-specific employer characteristics, which in other work are commonly assumed to be normally distributed.

Figure 5. Marginal distributions of estimated employer parameters



Note: Figure shows demeaned marginal distributions of estimated productivity (p), gender-specific amenity values (π_g), and gender wedges (z).

Table 6 reports employment-weighted pairwise correlations between gender-specific pay (ψ_g), gender-specific PageRanks (r_g), productivity (p), gender-specific amenity values (π_g), gender wedges (z), and vacancy cost shifters ($c_g^{v,0}$).³⁴ A few points are worth noting. First, within employers across genders, we find strong positive correlations between pay (0.900), ranks (0.651), amenities (0.662), and vacancy costs (0.672), indicating establishment-specific factors shared by men and women. Second, productivity is strongly positively correlated with ranks for men (0.847) but less so for women (0.586). This is because productivity is similarly positively correlated with pay for men (0.546) and women (0.582) but more positively correlated with amenities for men (0.556) than for women (0.247). In contrast, benchmark job-ladder models à la [Burdett and Mortensen \(1998\)](#) predict that productivity ranks are perfectly positively correlated with pay ranks. Third, the correlation between pay and ranks is more positive for men (0.414) than for women (0.349), while that between amenities and ranks is more positive for women (0.666) than for men (0.602). Fourth, amenities are similarly negatively correlated to pay for men (-0.331) and for women (-0.343), suggesting compensating differentials for both genders. Finally, gender wedges are correlated positively with ranks for men (0.376) but negatively for women (-0.281) and correlated positively with productivity (0.507). This is consistent with [Becker's \(1971\)](#) idea that taste-based discrimination cannot survive among employers with close-to-

³⁴In Appendix E.1, we report for robustness the relation between our estimates using PageRanks and analogous estimates using the poaching rank ([Moscarini and Postel-Vinay, 2008](#); [Bagger and Lentz, 2019](#)) and the net poaching rank ([Haltiwanger et al., 2018](#); [Moscarini and Postel-Vinay, 2018](#)).

zero economic profits.

Table 6. Correlation table for estimated employer parameters

	ψ_M	ψ_F	r_M	r_F	p	π_M	π_F	z	$c_M^{v,0}$	$c_F^{v,0}$
ψ_M	1.000									
ψ_F	0.900	1.000								
r_M	0.414	0.428	1.000							
r_F	0.277	0.349	0.651	1.000						
p	0.546	0.582	0.847	0.586	1.000					
π_M	-0.331	-0.245	0.602	0.420	0.556	1.000				
π_F	-0.341	-0.343	0.332	0.666	0.247	0.662	1.000			
z	0.363	0.238	0.376	-0.281	0.507	0.183	-0.403	1.000		
$c_M^{v,0}$	0.219	0.214	-0.170	-0.031	-0.174	-0.425	-0.226	-0.208	1.000	
$c_F^{v,0}$	0.361	0.334	0.025	-0.069	0.016	-0.340	-0.399	0.144	0.672	1.000

Note: Table reports employment-weighted pairwise correlations across employers between gender-specific pay (ψ_g), gender-specific PageRanks (r_g), productivity (p), gender-specific amenity values (π_g), gender wedges (z), and vacancy cost shifters ($c_g^{v,0}$).

6.4 Relating Estimates of Amenity Values to Observable Employer Characteristics

We estimate amenity values as the residuals that rationalize employer ranks given pay. To find out what economic factors these residuals capture, we relate our amenity estimates to a rich set of employer characteristics in the following specification:

$$\hat{\pi}_{g,j} = Z_{g,j}\eta_g + \iota_{g,j}, \quad (25)$$

where $\hat{\pi}_{g,j}$ is the estimated gender-specific amenity value for employer j and gender g , $Z_{g,j}$ is a vector of gender-specific employer covariates, and $\iota_{g,j}$ is an error term. We include as covariates in $Z_{g,j}$ 14 variables, which we construct using the RAIS data—see Appendix E.2 for details.

Table 7 shows results from estimating equation (25). Reassuringly, unmeasured income in the form of employer-provided food stamps loads similarly positively onto men’s and women’s amenity values. Both genders positively value amenities related to job flexibility and paid leave and negatively value attributes related to earnings fluctuations, workplace conflict (proxied by the share of unjust firings), and workplace risk (proxied by worker death rates). Compared with men, women put greater value on amenities such as hours flexibility and parental leave but greater disvalue on disamenities such as unpaid leave, earnings risk, and workplace risk.³⁵ Altogether, we explain around 32 percent

³⁵Parental leave may be costly to coworkers who increase working hours to make up for the temporary absence of a parent (Gallen, 2020; Brenøe et al., 2019; Ginja et al., 2020). While we find a large negative coefficient on parental leave incidence for men, men are empirically two orders of magnitude less likely than women to take parental leave.

of the variation in estimated amenities for men and 47 percent of the variation for women.

Table 7. Regressing estimates of amenity values on employer characteristics, by gender

	Men	Women
Indicator: employer provides food stamps	0.089*** (0.000)	0.083*** (0.000)
Share of workers with part-time contract	0.033*** (0.000)	0.096*** (0.000)
Share of workers with hours change since previous year	0.034*** (0.001)	0.123*** (0.001)
Share of workers with paid sick leave	0.175*** (0.001)	0.144*** (0.001)
Share of workers with parental leave	-4.969*** (0.036)	0.065*** (0.005)
Share of workers with unpaid leave	-0.085*** (0.004)	-0.125*** (0.005)
Share of workers with earnings cut since previous year	-0.165*** (0.001)	-0.219*** (0.001)
Share of workers with noncontractual-earnings fluctuations	-0.045*** (0.001)	-0.218*** (0.001)
Share of workers with work-related accident	-0.334*** (0.007)	-0.534*** (0.012)
Share of workers with commute-related accident	-0.792*** (0.026)	-0.311*** (0.044)
Share of worker separations due to firing for unjust reasons	-0.162*** (0.000)	-0.188*** (0.000)
Share of worker separations due to worker death	-0.627*** (0.003)	-0.786*** (0.004)
Industry FEs	✓	✓
Municipality FEs	✓	✓
Number of unique establishments	272,549	168,862
Observations	17,407,809	9,760,711
R^2	0.320	0.471

Note: Table reports estimated coefficients from regressing structurally estimated, gender-specific employer amenity values on observable gender-specific establishment characteristics. See equation (25). Details of covariates are presented in Appendix E.2. ***, **, and * denote significance at the 1 percent, 5 percent, and 10 percent levels, respectively.

6.5 Relating Estimates of Recruiting Costs to Observable Employer Characteristics

Our estimation allows us to identify gender-specific recruiting costs for each employer that rationalized observed hiring patterns in the data. The relative recruiting technology for women compared with men may differ owing to the gender specificity of human resource practices, referral networks, and local labor market tightness. To test for these postulated channels, we estimate the following specification:

$$\widehat{RRC}_j = Z_j \eta + \iota_j, \quad (26)$$

where $\widehat{RRC}_j \equiv \ln(\widehat{c}_F^{v,0} / \widehat{c}_M^{v,0})$ is the relative recruiting cost or log ratio of estimated female-to-male recruiting costs for employer j , Z_j is a vector of employer covariates, and ι_j is an error term. We include in Z_j nine covariates, which we construct using the RAIS data. See Appendix E.3 for details.

Table 8 presents the results from estimating equation (26). Focusing on column (3), which includes industry and municipality controls, we find that women's relative recruiting cost is lower at

employers that are larger and have a woman in the highest-paid position. Conversely, the relative cost of attracting women is higher at employers whose workforce consists of more college and high school graduates and older and higher-tenure employees. Comparing columns (1)–(3) highlights the explanatory power of industry and municipality effects in relative recruiting costs, with the most saturated specification in column (3) explaining around 54 percent of the estimated variation in relative recruiting costs. In summary, these findings are consistent with the postulated channels.

Table 8. Regressing estimates of women’s relative recruiting costs on employer characteristics

	(1)	(2)	(3)
Employer size	−0.045*** (0.000)	−0.047*** (0.000)	−0.043*** (0.000)
Indicator: highest-paid worker is a woman	−0.715*** (0.001)	−0.265*** (0.001)	−0.314*** (0.001)
Share of workers with college degree	0.181*** (0.002)	0.094*** (0.002)	0.344*** (0.002)
Share of workers with high school degree	0.350*** (0.002)	0.699*** (0.002)	0.509*** (0.002)
Mean worker age	0.335*** (0.003)	0.680*** (0.003)	0.523*** (0.003)
Mean worker tenure	−0.059*** (0.001)	0.151*** (0.000)	0.114*** (0.000)
Share of workers with commute-related accidents	2.052*** (0.142)	−2.637*** (0.113)	2.837*** (0.111)
Industry FEs		✓	✓
Municipality FEs			✓
Number of unique establishments	96,337	96,337	96,337
Observations	7,636,531	7,636,531	7,636,531
R ²	0.111	0.451	0.537

Note: Table reports estimated coefficients from regressing log ratio of structurally estimated recruiting costs of women to that of men on observable establishment characteristics—see equation (26). Details of covariates are presented in Appendix E.3. ***, **, and * denote significance at the 1 percent, 5 percent, and 10 percent levels, respectively.

6.6 Relating Estimates of Wedges to Observable Employer Characteristics

Gender wedges represent residuals that rationalize equilibrium pay of women with that of men within employers. Such wedges could capture taste-based discrimination (Becker, 1971) or employer-level comparative advantages across genders (Goldin, 1992; Rendall, 2018). To unpack this black box, we relate our estimates of gender wedges to employer characteristics in the following specification:

$$\hat{z}_j = Z_j \eta + \iota_j, \quad (27)$$

where \hat{z}_j is the estimated gender wedge for establishment j , Z_j is a vector of employer covariates, and ι_j is an error term. We include as covariates in Z_j 12 variables, which we construct using the RAIS data—see Appendix E.2 for details.

Table 9 contains the results from estimating equation (25). We group the independent variables into two categories. The first group pertains to proxies for employer-level comparative advantages across genders. It includes task content measures and proxies for risks of physical harm and strength requirements. Our results indicate mixed support for the employer-level-comparative-advantage story. On one hand, we find that gender wedges are positively related to nonroutine manual task intensity and the share of workers with work-related accidents, suggesting that women are less productive in physical and risky jobs.³⁶ On the other hand, we find that gender wedges are negatively related to routine manual task intensity and the share of worker separations due to worker death, suggesting that women are no less productive in physically demanding jobs.³⁷

The second group pertains to proxies for taste-based discrimination. It includes measures of workforce composition by gender and of financial independence. Our results lend some support to the interpretation of gender wedges as being related to taste-based discrimination. Gender wedges are negatively related to the female employment share, consistent with more female-friendly workplaces' attracting and retaining more women. Establishments with a woman in the highest-paid position have lower average gender wedges, suggesting that female mentors may affect the female friendliness of a workplace.³⁸ Finally, small firms with little financial dependence have a higher gender wedge, possibly because discrimination is more likely to survive under lower external accountability.

These results are robust across the specifications in columns (1)–(3) with cumulatively added controls for employers' industry and municipality, which may themselves be correlated with gender-based comparative advantage or taste-based discrimination.

Although we do not definitively pin down the factors behind the gender wedges, our analysis sheds light on proxies related to potential economic explanations. Altogether, we explain around 76 percent of the variation in estimated gender wedges. Most of this variation is due to variables related to gender-based comparative advantage and taste-based discrimination rather than due to the inclusion of industry and municipality controls. While our findings suggest the economic explanations for the estimated gender wedges, they also leave room for other explanations.

³⁶*Nonroutine manual tasks* include jobs that are physical but involve a wide set of tasks and are not rule based, such as “janitor, home health aide, and personal care aide” (Siu and Jaimovich, 2015).

³⁷*Routine manual tasks* include “jobs that are both rule based and emphasize physical [...] tasks,” such as “factory workers who operate welding, fitting, and metal press machines [and] forklift operators” (Siu and Jaimovich, 2015).

³⁸This analysis is potentially subject to a simultaneity issue—namely, that higher gender wedges might lead employers to hire fewer women. For robustness, Table E.3 in Appendix E.6 presents results of the same regression with the variables “female employment share” and “indicator: highest-paid worker is a woman” omitted, with largely similar results.

Table 9. Regressing estimates of gender wedges on employer characteristics

	(1)	(2)	(3)
Routine manual task intensity	-0.107*** (0.000)	-0.059*** (0.001)	-0.057*** (0.001)
Nonroutine manual task intensity	0.278*** (0.001)	0.176*** (0.001)	0.155*** (0.001)
Routine cognitive task intensity	-0.013*** (0.000)	-0.005*** (0.001)	0.003*** (0.001)
Nonroutine cognitive interpersonal task intensity	-0.123*** (0.001)	-0.029*** (0.001)	-0.030*** (0.001)
Nonroutine cognitive analytical task intensity	0.089*** (0.001)	0.055*** (0.001)	0.034*** (0.001)
Share of worker separations due to worker death	-0.753*** (0.005)	-0.471*** (0.005)	-0.395*** (0.006)
Share of workers with work-related accidents	2.229*** (0.021)	1.500*** (0.021)	0.295*** (0.020)
Female employment share	-4.206*** (0.001)	-3.645*** (0.001)	-3.835*** (0.001)
Indicator: highest-paid worker is a woman	-0.239*** (0.001)	-0.166*** (0.001)	-0.121*** (0.001)
Indicator: no major financial stakeholders	0.048*** (0.001)	0.031*** (0.001)	0.034*** (0.001)
Industry FEs		✓	✓
Municipality FEs			✓
Number of unique establishments	96,065	96,065	96,065
Observations	17,287,101	17,287,101	17,287,101
R^2	0.693	0.730	0.764

Note: Table reports estimated coefficients from regressing structurally estimated gender wedges on observable establishment characteristics. See equation (27). Details of covariates are presented in Appendix E.5. ***, **, and * denote significance at the 1 percent, 5 percent, and 10 percent levels, respectively.

6.7 Model Fit

We solve for the equilibrium of the model given the above parameter estimates. We first use the model to match exactly, employer by employer, the empirical offer distribution of pay and amenities F_g^r . We then ask the model to predict the cross-sectional equilibrium pay and amenity distribution, which we compare with the data.³⁹ Table 10 shows the model fit vis-à-vis a set of salient empirical moments. We find that the model fits the data well. Compared with the data, the model somewhat understates the magnitude of the gender pay gap and the variance of employer pay. The model closely matches the empirical variance of the gender pay gap, empirical job-to-job transition rates for both genders, and the correlation between men's and women's pay within employers.

7 Equilibrium Counterfactuals

7.1 Sources of the Gender Pay Gap

With the estimated model in hand, we recompute equilibria while shutting down various gender differences. We consider four counterfactuals. First, we set the amenity cost shifters of women equal to

³⁹The numerical solution algorithm we use is described in Appendix E.7.

Table 10. Model fit

Moment	Description	Data	Model
$\mathbb{E}[\psi_M - \psi_F]$	Gender pay gap	0.084	0.074
$\mathbb{E}[\psi_F g = M] - \mathbb{E}[\psi_F g = F]$	Gender pay gap between employers	0.074	0.055
$\mathbb{E}[\psi_M - \psi_F g = F]$	Gender pay gap within employers	0.009	0.018
$Var(\psi_M)$	Variance of men's pay	0.051	0.040
$Var(\psi_F)$	Variance of women's pay	0.046	0.032
$Var(\psi_M - \psi_F)$	Variance of gender pay gap	0.010	0.009
$\mathbb{E}[\lambda_M^e (1 - F_M(x)) + \lambda_M^G]$	Job-to-job transition rate for men	0.016	0.015
$\mathbb{E}[\lambda_F^e (1 - F_F(x)) + \lambda_F^G]$	Job-to-job transition rate for women	0.012	0.012
$Corr(\psi_M, \psi_F)$	Correlation between men's and women's pay	0.926	0.932

Note: Table reports model fit with respect to data-based and model-based moments relating to gender-specific employer pay (ψ_g) and monthly job-to-job transition rates ($\lambda_g^e (1 - F_g(x)) + \lambda_g^G$) for $g = M, F$.

those of men, $c_F^{\pi,0} = c_M^{\pi,0}$.⁴⁰ Second, we shut down gender wedges, $z = 0$. Third, we shut down only differences in vacancy creation cost shifters $c_g^{v,0}$ across genders. Fourth and finally, we simultaneously shut down all of the above gender differences.⁴¹ It is important to note that these counterfactuals not only remove gender differences in model parameters but also change population means of these parameters by setting women's parameters equal to those of men. We think of this as a natural benchmark in a world in which the labor market for men is relatively undistorted. Table 11 describes our baseline economy (column 0) and results from the four counterfactuals (columns 1–4).

Counterfactual 1: Shutting down gender differences in amenities. Removing gender differences in amenities—see column (1) of Table 11—accounts for around 1.3 log points (18 percent) of the gender gap. Worker welfare increases by 0.5 percent in payroll-equivalent units. The increase in pay for women (Panel A of Figure 6) is offset partly by a corresponding decrease in mean amenity values (Panel B). Output increases slightly because high-productivity firms increase their amenities, thus attracting more women. As a consequence, pay inequality for women increases. However, employer welfare decreases: profits decline because firms now pay higher wages, and gender wedges increase because more women work at high-wedge firms. Finally, women's employer rank-pay rank relationship becomes steeper (Panel C), becoming more like that of men documented in Section 3.2.

Counterfactual 2: Shutting down gender wedges. Removing gender wedges—see column (2) of Table 11—accounts for around 5.4 log points (73 percent) of the gender gap. Most of the decrease in

⁴⁰This counterfactual is equivalent to one that equates gender preferences over an amenity vector, which we describe in Appendix C.4. For robustness, we simulated counterfactuals under a range of amenity cost elasticities. In Appendix F, we show a variant of the counterfactuals with exogenous amenities, or $\eta_\pi = \infty$. Our insights remain substantially unchanged.

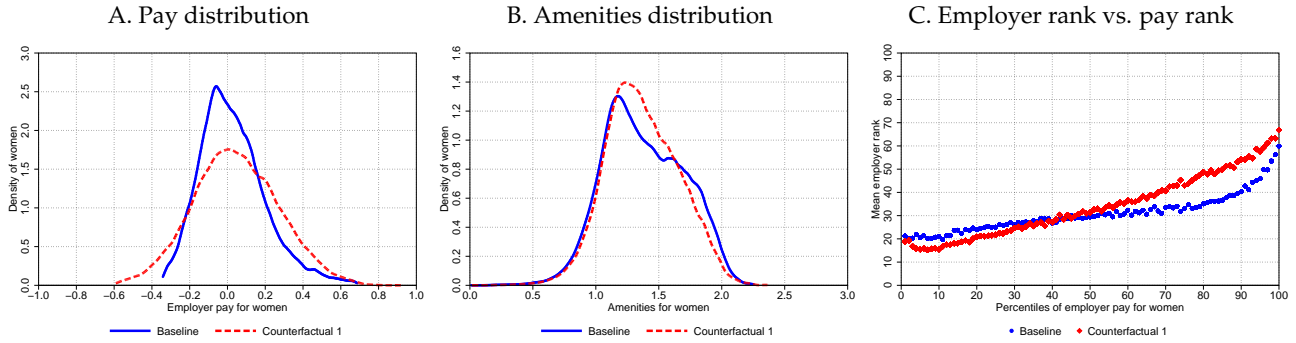
⁴¹In all counterfactuals but the last, we keep unchanged single-gender employers, which make up a significant share of all employers and mediate the equilibrium response of other firms to counterfactual changes. For the gender-neutral economy, previously all-male firms hire women in equal proportions, while previously all-female firms cease to exist.

Table 11. Equilibrium decomposition of gender pay gap

	Baseline	Counterfactuals			
	(0)	(1)	(2)	(3)	(4)
Gender differences in...					
...amenities	✓		✓	✓	
...employer wedges	✓	✓		✓	
...vacancy posting costs	✓	✓	✓		
Gender pay gap...	0.074	0.061	0.020	0.018	0.000
...between employers	0.055	0.056	0.047	0.016	0.000
...within employers	0.018	0.005	-0.026	0.002	0.000
Output	1.000	1.001	1.012	1.033	1.035
Worker welfare...	0.000	0.004	0.015	-0.004	0.027
...from payroll for women	0.000	0.010	0.026	0.019	0.029
...from amenity value for women	0.000	-0.006	-0.010	-0.022	-0.002
Payroll-equivalent welfare change	0.000	0.005	0.019	-0.004	0.033
Employer welfare...	1.000	0.997	1.011	0.986	1.039
...from profits	1.004	1.002	1.011	1.039	1.039
...from wedges	-0.004	-0.006	0.000	-0.053	0.000
Total employment for men	0.757	0.757	0.757	0.757	0.757
Total employment for women	0.760	0.759	0.762	0.760	0.757

Note: Table reports simulation results from model-based counterfactuals. Baseline results (column 0) are compared against counterfactuals without gender differences in amenities (column 1), in employer gender wedges (column 2), in vacancy posting costs (column 3), and without any gender differences (column 4).

Figure 6. Effects of counterfactual 1: removing gender differences in amenities

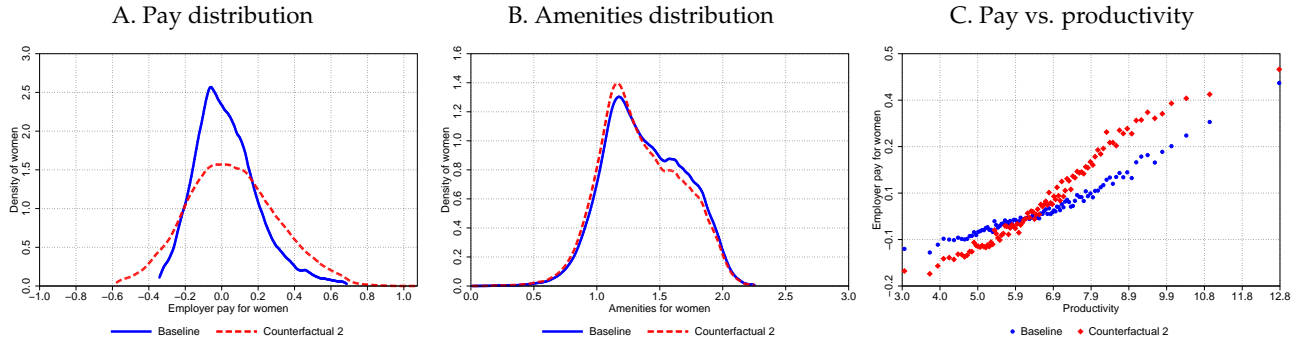


Note: Figure shows women's pay distributions (Panel A), amenities distributions (Panel B), and employer rank-pay rank relationship (Panel C) for counterfactual corresponding to column (1) of Table 11.

the gender gap is within firms, but the between-employer gap also decreases because of equilibrium reallocation of women. The mean and dispersion of women's wages increase (Panel A of Figure 7). Since gender wedges were positively correlated with productivity, removing them increases women's dispersion in composite productivity and hence pay. Output increases by 1.2 percent, worker welfare by 1.9 percent, and profits by 0.7 percent. Output increases owing to women's relocating to high-productivity, formerly high-wedge firms. The increase in worker welfare is accounted for by a large rise in women's payroll and offset partly by a decline in amenity values due to worker relocation

(Panel B). That employer profits also increase reflects the fact that employers with a nonzero gender wedge were not maximizing monetary profits. Another result is a steepening of the productivity-pay relationship for women (Panel C) that resembles an increase in women’s “bargaining power” (Card et al., 2016).

Figure 7. Effects of counterfactual 2: removing gender differences in gender wedges

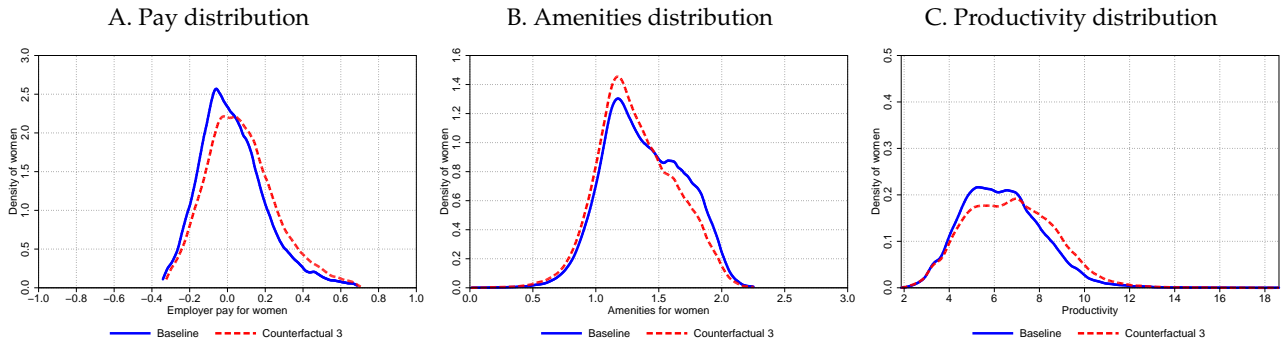


Note: Figure shows women’s pay distributions (Panel A), amenities distributions (Panel B), and employer pay-productivity relationship (Panel C) for counterfactual corresponding to column (2) of Table 11.

Counterfactual 3: Shutting down gender differences in vacancy costs. Removing gender differences in vacancy costs—see column (3) of Table 11—accounts for 5.6 log points (76 percent) of the gender gap. Again, most of this decline is explained by a decline in the between-employer pay gap. More women become employed at high-paying firms (Panel A of Figure 8). There is a significant increase in output of around 3.3 percent because of the relocation of women to more productive firms (Panel C). However, the impact on worker welfare is net negative, around -0.4 percent. The reason for this is that especially high-wedge, high-productivity, but low-amenity employers increase their employment of women (Panel C). This also means that in spite of an increase of around 3.5 percent in monetary profits, employers are worse off on average. A key takeaway from this simulation is that gender differences in employer allocation are not necessarily inefficient.

Counterfactual 4: Moving to a gender-neutral economy. By construction, moving to a gender-neutral economy—see column (4) of Table 11—eliminates the gender gap. Put differently, women’s mean pay increases by 7.4 percent. Interestingly, this is also associated with large gains in output (3.5 percent), worker welfare (3.3 percent), and employer welfare (3.9 percent). Output increases because women relocate to higher-productivity employers. Most of the increase in worker welfare is due to an increase in pay, not amenities. Additionally, employer welfare increases because of a combination of higher profits and lower gender wedges.

Figure 8. Effects of counterfactual 3: removing gender differences in vacancy posting costs



Note: Figure shows women’s pay distributions (Panel A), amenities distributions (Panel B), and productivity distributions (Panel C) for counterfactual corresponding to column (3) of Table 11.

Interaction effects and the distinction between output and welfare. One important insight from our simulations is that the different structural gender differences interact nonlinearly. While simultaneously removing all gender differences leads to large output and welfare gains, addressing only one at a time may actually result in welfare losses. A second important insight is that output and welfare are fundamentally different and that the two can move in opposite directions.

7.2 The (Unintended) Equilibrium Consequences of Equal-Treatment Policies

That gender gaps in pay, employment, and amenities across firms are associated with sizable output and welfare losses suggests three interesting thought experiments. First, what would be the effects of an equal-pay policy that requires men and women of the same ability to be equally paid within employers? Second, what would be the effects of an equal-hiring policy requiring employers to advertise the same number of job openings for men and women? Third, what would be the effects of an equal-amenities policy that requires employers to offer the same amenity values to men and women? Table 12 presents the results of simulating these counterfactual policies.⁴² Our overarching insight is that the (unintended) equilibrium consequences of forcing equal pay, amenities, or hiring on employers are quite different from moving to a gender-neutral economy.

Effects of an equal-pay policy. By construction, an equal-pay policy forcing employers to set the same wage for men and women eliminates the within-employer gender pay gap. More surprisingly, the policy increases the between-employer gender pay gap by 0.2 log points. This happens because

⁴²The numerical solution algorithm used to solve for the equilibrium under the simulated policies departs significantly from that used for the previous results and is detailed in Appendix F.1. For robustness, we repeat our simulations under the assumption of exogenous amenities, results of which can be found in Table F.2 of Appendix F.

firms with a positive gender wedge find it especially costly to pay a single wage to workers of both genders, leading them to reduce hiring of women in equilibrium. Because gender wedges are higher at high-productivity firms, the equal-pay policy reduces women's employment at relatively high-paying firms, thereby increasing the between-employer gap. The policy also has subtle redistributive effects. The policy increases women's pay by 1.6 log points, while marginally decreasing men's pay. Employers compensate part of the pay changes for each gender with opposite changes in amenities.⁴³

On aggregate, the equal-pay policy has little effect on output and welfare. Worker welfare increases, as the rise in women's payroll exceeds the drop in women's amenities plus men's loss of income net of amenities. Employer welfare decreases, owing mostly to lower profits.

Effects of an equal-hiring policy. An equal-hiring policy forcing employers to post the same number of vacancies for men and women largely equalizes the distributions of men and women across employers. As a result, the between-employer pay gap drops from 5.5 to 0.2 log points. However, the policy has a number of adverse consequences on the economy, because it distorts gender-specific allocative efficiency. Notably, the within-employer gap increases. This happens because firms with positive gender wedges and relatively higher costs of producing amenities for women find it especially costly to balance hiring. These firms adjust to the equal-hiring policy by lowering women's wages and by decreasing overall hiring, which decreases men's employment by 0.3 percentage points and women's employment by 1.2 percentage points. Consequently, there is a slight reduction in overall payroll and a large reduction in amenities, with about two-thirds of the latter concentrated among women.

Overall, this policy results in a large decline in worker welfare by 4.7 log points in payroll-equivalent terms, a decrease in output of 2.1 log points, and a reduction in employer welfare by 5.1 log points.

Effects of an equal-amenities policy. Perhaps surprisingly, an equal-amenities policy forcing employers to offer the same amenity value to men and women reduces the within-employer pay gap but, at the same time, dramatically increases the between-employer pay gap, thereby increasing the overall gender pay gap by 5.1 log points. This happens because firms on average adjust amenities upward for women but downward for men. Since this increases the marginal cost of amenity production above one for women while decreasing below one that for men, firms lower women's wages and increase men's wages in equilibrium. On top of this adjustment, there is large reallocation of

⁴³In this sense, our findings are comparable to empirical estimates of the effects of equal-pay legislation (Manning, 1996) and related, but different, pay transparency laws (Bennedsen et al., 2020).

employment of women from high-pay but low-amenities firms to low-pay but high-amenities firms. The equal-amenities policy skews employment of these low-paying firms further toward women and away from men, thereby increasing the between-employer gender pay gap.

The changes in the wage structure due to the equal-amenities policy have muted effects on aggregate output, worker welfare, and employer welfare due to offsetting pay and amenities in equilibrium.

Table 12. Effects of simulated equal-pay, equal-hiring, and equal-amenities policies

	Baseline	Simulated policies		
		Equal-pay	Equal-hiring	Equal-amenities
	(0)	(1)	(2)	(3)
Mean pay for men	1.000	1.000	1.006	1.019
Mean pay for women	1.000	1.016	1.031	0.968
Gender pay gap...	0.074	0.057	0.049	0.125
...between employers	0.055	0.057	0.002	0.138
...within employers	0.018	0.000	0.047	-0.013
Output	1.000	1.000	0.979	1.000
Worker welfare...	0.000	0.001	-0.038	0.000
...from total payroll...	0.000	0.002	-0.002	0.005
...for men	0.000	-0.001	0.001	0.009
...for women	0.000	0.003	-0.003	-0.004
...from total amenity value...	0.000	-0.001	-0.036	-0.004
...for men	0.000	0.001	-0.011	-0.008
...for women	0.000	-0.002	-0.025	0.004
Payroll-equivalent welfare change	0.000	0.001	-0.047	0.000
Employer welfare...	1.000	0.999	0.949	0.999
...from profits	1.004	1.003	0.978	1.002
...from wedges	-0.004	-0.004	-0.029	-0.003
Total employment for men	0.757	0.757	0.754	0.757
Total employment for women	0.760	0.760	0.748	0.760

Note: Table reports results from simulating counterfactual policies. Baseline results (column 0) are compared against the economy under an equal-pay policy (column 1), an equal-hiring policy (column 2) and an equal-amenities policy (column 3).

8 Conclusion

Our analysis sheds light on the micro sources and macro consequences of the gender pay gap. We document that a large share of the gender pay gap in Brazil is accounted for by women working at lower-paying employers, compared with men. At the same time, workers seem to have gender-specific preferences over nonpay employer characteristics. We develop an empirical equilibrium search model that can rationalize the gender differences in pay and revealed-preference ranks in the

data. We use the estimated model to simulate a series of model experiments, including a structural decomposition of the gender pay gap and policy counterfactuals.

Does the fact that women are paid less than men reflect output or welfare losses? The answer depends on the microstructure of the labor market. In our specific context, we find that some ways of closing the gender gap increase utility for women, while others leave them worse off. This does not mean that gender gaps should be ignored. We find sizable output and welfare gains from moving to a gender-neutral economy. Nevertheless, achieving gender equality may not be an easy task. Our analysis suggests that equal-treatment policies fail to close the gender pay gap in equilibrium.

Our work opens up several avenues for future research. We have been agnostic about some of the factors underlying gender differences in labor market parameters. Additional evidence could shed light on these factors and help address them with policies. The applicability of our framework is also not limited to gender. The theoretical framework and estimation routine we develop can be used to assess other population subgroups by race, cohort, or educational background. By allowing for separate job ladders for men and women, our work takes but a first step in the direction of integrating richer heterogeneity into equilibrium models of the labor market.

References

- Abowd, John M., Francis Kramarz, and David N. Margolis**, "High Wage Workers and High Wage Firms," *Econometrica*, 1999, 67 (2), 251–333.
- Abraham, Lisa and Alison Stein**, "Words Matter: Experimental Evidence from Job Applications," Working Paper, 2020.
- Acemoglu, Daron and David Autor**, "Skills, Tasks and Technologies: Implications for Employment and Earnings," in O. Ashenfelter and D. Card, eds., *Handbook of Labor Economics*, Vol. 4, Amsterdam: Elsevier-North, 2011, chapter 12, pp. 1043–1171.
- Albanesi, Stefania**, "Changing Business Cycles: The Role of Women's Employment," Working Paper, 2020.
- and **Ayşegül Şahin**, "The Gender Unemployment Gap," *Review of Economic Dynamics*, 2018, 30, 47–67.
- and **Claudia Olivetti**, "Home Production, Market Production and the Gender Wage Gap: Incentives and Expectations," *Review of Economic Dynamics*, 2009, 12 (1), 80–107.
- and – , "Gender Roles and Medical Progress," *Journal of Political Economy*, 2016, 124 (3), 650–695.
- Albrecht, James, Carlos Carrillo-Tudela, and Susan Vroman**, "On-the-Job Search with Match-Specific Amenities," *Economics Letters*, 2018, 162, 15–17.
- Alvarez, Jorge, Felipe Benguria, Niklas Engbom, and Christian Moser**, "Firms and the Decline in Earnings Inequality in Brazil," *American Economic Journal: Macroeconomics*, 2018, 10 (1), 149–189.
- Amano-Patiño, Noriko, Tatiana Baron, and Pengpeng Xiao**, "Equilibrium Wage-Setting and the Life-Cycle Gender Pay Gap," Working Paper, 2019.
- Andrews, Martyn J., Leonard Gill, Thorsten Schank, and Richard Upward**, "High Wage Workers and Low Wage Firms: Negative Assortative Matching or Limited Mobility Bias?," *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 2008, 171 (3), 673–697.
- , – , – , and – , "High Wage Workers Match with High Wage Firms: Clear Evidence of the Effects of Limited Mobility Bias," *Economics Letters*, 2012, 117 (3), 824–827.

- Arrow, Kenneth J.**, "Some Models of Racial Discrimination in the Labor Market," Technical Report, Rand Corporation, Santa Monica, CA 1971.
- Autor, David and David Dorn**, "This Job Is 'Getting Old': Measuring Changes in Job Opportunities Using Occupational Age Structure," *American Economic Review*, 2009, 99 (2), 45–51.
- Bagger, Jesper and Rasmus Lentz**, "An Empirical Model of Wage Dispersion with Sorting," *The Review of Economic Studies*, 2019, 86 (1), 153–190.
- , **Bent Jesper Christensen, and Dale T. Mortensen**, "Productivity and Wage Dispersion: Heterogeneity or Misallocation?," Working Paper, 2014.
- Becker, Gary S.**, *The Economics of Discrimination*, University of Chicago Press, 1971.
- Bennedsen, Morten, Elena Simintzi, Margarita Tsoutsoura, and Daniel Wolfenzon**, "Do Firms Respond to Gender Pay Gap Transparency?," NBER Working Paper No. 25435, 2020.
- Black, Dan A.**, "Discrimination in an Equilibrium Search Model," *Journal of Labor Economics*, 1995, 13 (2), 309–334.
- Bonhomme, Stéphane and Grégory Jolivet**, "The Pervasive Absence of Compensating Differentials," *Journal of Applied Econometrics*, 2009, 24 (5), 763–795.
- , **Kerstin Holzheu, Thibaut Lamadon, Elena Manresa, Magne Mogstad, and Bradley Setzler**, "How Much Should We Trust Estimates of Firm Effects and Worker Sorting?," NBER Working Paper 27368, 2020.
- , **Thibaut Lamadon, and Elena Manresa**, "A Distributional Framework for Matched Employer Employee Data," *Econometrica*, 2019, 87 (3), 699–739.
- Bontemps, Christian, Jean-Marc Robin, and Gerard J. van den Berg**, "An Empirical Equilibrium Job Search Model with Search on the Job and Heterogeneous Workers and Firms," *International Economic Review*, 1999, 40 (4), 1039–1074.
- , – , and – , "Equilibrium Search with Continuous Productivity Dispersion: Theory and Nonparametric Estimation," *International Economic Review*, 2000, 41 (2), 305–358.
- Borovičková, Katarína and Robert Shimer**, "High Wage Workers Work for High Wage Firms," Working Paper, 2020.
- Bowlus, Audra J.**, "A Search Interpretation of Male-Female Wage Differentials," *Journal of Labor Economics*, 1997, 15 (4), 625–657.
- Brenøe, Anne A., Serena P. Canaan, Nikolaj A. Harmon, and Heather N. Royer**, "Is Parental Leave Costly for Firms and Coworkers?," IZA Discussion Paper 12870, 2019.
- Burdett, Kenneth and Dale T. Mortensen**, "Wage Differentials, Employer Size, and Unemployment," *International Economic Review*, 1998, 39 (2), 257–273.
- Card, David, Ana Rute Cardoso, and Patrick Kline**, "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women," *Quarterly Journal of Economics*, 2016, 131 (2), 633–686.
- , – , **Jörg Heining, and Patrick Kline**, "Firms and Labor Market Inequality: Evidence and Some Theory," *Journal of Labor Economics*, 2018, 36 (S1), S13–S70.
- , **Jörg Heining, and Patrick Kline**, "Workplace Heterogeneity and the Rise of West German Wage Inequality," *Quarterly Journal of Economics*, 2013, 128 (3), 967–1015.
- Cortés, Patricia, Jessica Pan, Laura Pilossoph, and Basit Zafar**, "Gender Differences in Job Search and the Earnings Gap: Evidence from Business Majors," Working Paper, 2020.
- Coudin, Elise, Sophie Maillard, and Maxime Tô**, "Family, Firms and the Gender Wage Gap in France," Institute for Financial Studies Working Paper W18/01, 2018.
- Dey, Matthew S. and Christopher J. Flinn**, "An Equilibrium Model of Health Insurance Provision and Wage Determination," *Econometrica*, 2005, 73 (2), 571–627.
- Dix-Carneiro, Rafael, Pinelopi Goldberg, Costas Meghir, and Gabriel Ulyssea**, "Trade and Informality in the Presence of Labor Market Frictions and Regulations," Working Paper, 2019.
- Engbom, Niklas and Christian Moser**, "Earnings Inequality and the Minimum Wage: Evidence from Brazil," Federal Reserve Bank of Minneapolis OIGI Working Paper 7, 2018.
- Erosa, Andrés, Luisa Fuster, Gueorgui Kambourov, and Richard Rogerson**, "Hours, Occupations, and Gender Differences in Labor Market Outcomes," Working Paper, 2019.
- Fernández, Raquel, Alessandra Fogli, and Claudia Olivetti**, "Mothers and Sons: Preference Forma-

- tion and Female Labor Force Dynamics," *The Quarterly Journal of Economics*, 2004, 119 (4), 1249–1299.
- Flabbi, Luca**, "Gender Discrimination Estimation in a Search Model with Matching and Bargaining," *International Economic Review*, 2010, 51 (3), 745–783.
- Flemming, Jean**, "Costly Commuting and the Job Ladder," Working Paper, 2020.
- Fogli, Alessandra and Laura Veldkamp**, "Nature or Nurture? Learning and the Geography of Female Labor Force Participation," *Econometrica*, 2011, 79 (4), 1103–1138.
- Folke, Olle and Johanna Rickne**, "Sexual Harassment and Gender Inequality in the Labor Market," Working Paper, 2020.
- Fukui, Masao, Emi Nakamura, and Jón Steinsson**, "Women, Wealth Effects, and Slow Recoveries," Working Paper, 2020.
- Gallen, Yana**, "Is Parental Leave Costly for Firms and Coworkers?," NBER Working Paper No. 26622, 2020.
- Gerard, François, Lorenzo Lagos, Edson Severnini, and David Card**, "Assortative Matching or Exclusionary Hiring? The Impact of Firm Policies on Racial Wage Differences in Brazil," NBER Working Paper No. 25176, 2018.
- Ginja, Rita, Arizo Karimi, and Pengpeng Xiao**, "Family Leave Programs: Employer Responses and the Gender Wage Gap," Working Paper, 2020.
- Goldin, Claudia**, *Understanding the Gender Gap: An Economic History of American Women*, Oxford University Press, 1992.
- , "A Grand Gender Convergence: Its Last Chapter," *American Economic Review*, 2014, 104 (4), 1091–1119.
- **and Cecilia Rouse**, "Orchestrating Impartiality: The Impact of "Blind" Auditions on Female Musicians," *American Economic Review*, September 2000, 90 (4), 715–741.
- Hall, Robert E. and Andreas I. Mueller**, "Wage Dispersion and Search Behavior: The Importance of Nonwage Job Values," *Journal of Political Economy*, 2018, 126 (4), 1594–1637.
- Haltiwanger, John C., Henry R. Hyatt, Lisa B. Kahn, and Erika McEntarfer**, "Cyclical Job Ladders by Firm Size and Firm Wage," *American Economic Journal: Macroeconomics*, 2018, 10 (2), 52–85.
- Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante**, "The Macroeconomics of the Quiet Revolution: Understanding the Implications of the Rise in Women's Participation for Economic Growth and Inequality," *Research in Economics*, 2017, 71 (3), 521–539.
- Heckman, James J., Lance J. Lochner, and Petra E. Todd**, "Earnings Functions, Rates of Return and Treatment Effects: The Mincer Equation and Beyond," in Erik Hanushek and F. Welch, eds., *Handbook of the Economics of Education*, Vol. 1, Amsterdam: Elsevier, 2006, chapter 7, pp. 307–458.
- Heise, Sebastian and Tommaso Porzio**, "Spatial Wage Gaps in Frictional Labor Markets," Federal Reserve Bank of Minneapolis OIGI Working Paper 29, 2019.
- Hellerstein, Judith K., Melissa McInerney, and David Neumark**, "Neighbors and Coworkers: The Importance of Residential Labor Market Networks," *Journal of Labor Economics*, 2011, 29 (4), 659–695.
- Hsieh, Chang-Tai and Peter J. Klenow**, "Misallocation and Manufacturing TFP in China and India," *The Quarterly Journal of Economics*, 2009, 124 (4), 1403–1448.
- , **Erik Hurst, Charles I. Jones, and Peter J. Klenow**, "The Allocation of Talent and U.S. Economic Growth," *Econometrica*, 2019, 87 (5), 1439–1474.
- Hwang, Hae-shin, Dale T. Mortensen, and W. Robert Reed**, "Hedonic Wages and Labor Market Search," *Journal of Labor Economics*, 1998, 16 (4), 815–847.
- Jarosch, Gregor**, "Searching for Job Security and the Consequences of Job Loss," Working Paper, 2015.
- Jolivet, Grégory, Fabien Postel-Vinay, and Jean-Marc Robin**, "The Empirical Content of the Job Search Model: Labor Mobility and Wage Distributions in Europe and the US," *European Economic Review*, 2006, 50 (4), 877–907.
- Jung, Philip and Moritz Kuhn**, "Earnings Losses and Labor Mobility over the Life Cycle," Working Paper, 2017.
- Kleven, Henrik, Camille Landais, and Jakob Egholt Søggaard**, "Children and Gender Inequality: Evidence from Denmark," *American Economic Journal: Applied Economics*, October 2019, 11 (4), 181–209.
- Kleven, Henrik Jacobsen, Camille Landais, and Jakob Egholt Søggaard**, "Children and Gender In-

- equality: Evidence from Denmark," Working Paper, 2016.
- Kline, Patrick, Raffaele Saggio, and Mikkel Sølvsten**, "Leave-Out Estimation of Variance Components," NBER Working Paper 26244, 2019.
- Lamadon, Thibaut, Magne Mogstad, and Bradley Setzler**, "Imperfect Competition, Compensating Differentials and Rent Sharing in the U.S. Labor Market," Working Paper No. 2019-84, Becker Friedman Institute for Economics, University of Chicago, 2019.
- Lang, Kevin and Sumon Majumdar**, "The Pricing of Job Characteristics When Markets Do Not Clear: Theory and Policy Implications," *International Economic Review*, 2004, 45 (4), 1111–1128.
- Lavetti, Kurt and Ian M. Schmutte**, "Estimating Compensating Wage Differentials with Endogenous Job Mobility," Working Paper 29, Labor Dynamics Institute, Cornell University 2018.
- Lentz, Rasmus and Dale T. Mortensen**, "An Empirical Model of Growth Through Product Innovation," *Econometrica*, 2008, 76 (6), 1317–1373.
- Luo, Mi and Simon Mongey**, "Assets and Job Choice: Student Debt, Wages and Job Satisfaction," NBER Working Paper 25801, 2019.
- Manning, Alan**, "The Equal Pay Act as an Experiment to Test Theories of the Labour Market," *Economica*, 1996, 63 (250), 191–212.
- Meghir, Costas, Renata Narita, and Jean-Marc Robin**, "Wages and Informality in Developing Countries," *American Economic Review*, 2015, 105 (4), 1509–1546.
- Mincer, Jacob A.**, *Schooling, Experience, and Earnings*, New York: NBER Press, 1974.
- Mortensen, Dale T.**, *Wage Dispersion: Why Are Similar Workers Paid Differently?*, The MIT Press, 2003.
- Moscarini, Giuseppe and Fabien Postel-Vinay**, "The Timing of Labor Market Expansions: New Facts and a New Hypothesis," in Rogoff K. Acemoglu, D. and M. Woodford, eds., *NBER Macroeconomics Annual*, Vol. 23, University of Chicago Press, 2008, chapter 1, pp. 1–52.
- and – , "Stochastic Search Equilibrium," *Review of Economic Studies*, 2013, 80 (4), 1545–1581.
- and – , "The Cyclical Job Ladder," *Annual Review of Economics*, 2018, 10, 165–188.
- Ngai, L. Rachel and Barbara Petrongolo**, "Gender Gaps and the Rise of the Service Economy," *American Economic Journal: Macroeconomics*, 2017, 9 (4), 1–44.
- Page, Lawrence, Sergey Brin, Rajeev Motwani, and Terry Winograd**, "The PageRank Citation Ranking: Bringing Order to the Web," Working Paper, 1998.
- Peters, Michael and Conor Walsh**, "Declining Dynamism, Increasing Markups and Missing Growth: The Role of the Labor Force," Working Paper, 2019.
- Rendall, Michelle**, "Brain versus Brawn: The Realization of Women's Comparative Advantage," Working Paper No. 491, Institute For Empirical Research in Economics, University of Zurich, 2018.
- Restuccia, Diego and Richard Rogerson**, "Policy Distortions and Aggregate Productivity with Heterogeneous Plants," *Review of Economic Dynamics*, 2008, 11 (4), 707–720.
- Rosen, Sherwin**, "The Theory of Equalizing Differences," in O. Ashenfelter and R. Layard, eds., *Handbook of Labor Economics*, Vol. 1, Amsterdam: Elsevier-North, 1986, chapter 12, pp. 641–692.
- Siu, Henry and Nir Jaimovich**, "Jobless Recoveries," Report, Third Way, Washington DC 2015.
- Song, Jae, David J. Price, Fatih Guvenen, Nicholas Bloom, and Till Von Wachter**, "Firming Up Inequality," *Quarterly Journal of Economics*, 2018, 134 (1), 1–50.
- Sorkin, Isaac**, "Ranking Firms Using Revealed Preference and Other Essays about Labor Markets," Doctoral Dissertation, University of Michigan, 2015.
- , "The Role of Firms in Gender Earnings Inequality: Evidence from the United States," *American Economic Review*, 2017, 107 (5), 384–387.
- , "Ranking Firms Using Revealed Preference," *The Quarterly Journal of Economics*, 2018, 133 (3), 1331–1393.
- Sullivan, Paul and Ted To**, "Search and Nonwage Job Characteristics," *Journal of Human Resources*, 2014, 49 (2), 472–507.
- Taber, Christopher and Rune Vejlin**, "Estimation of a Roy/Search/Compensating Differential Model of the Labor Market," *Econometrica*, 2020, 88 (3), 1031–1069.
- Ulysea, Gabriel**, "Firms, Informality, and Development: Theory and Evidence from Brazil," *American Economic Review*, 2018, 108 (8), 2015–2047.
- Xiao, Pengpeng**, "Wage and Employment Discrimination by Gender in Labor Market Equilibrium,"

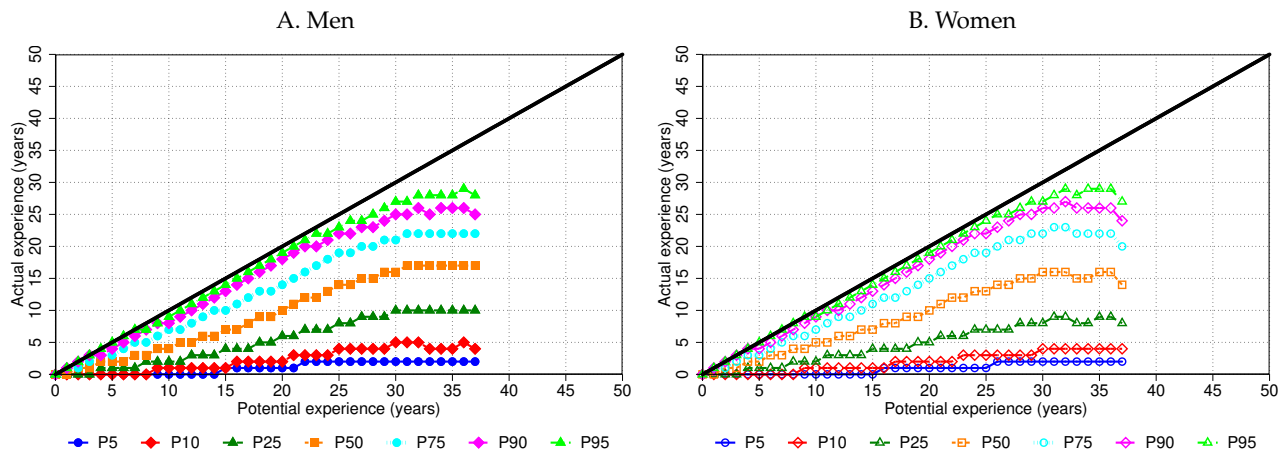
Working Paper, 2020.

Online Appendix—Not for Publication

A Data Appendix

A.1 Comparison of Actual versus Potential Experience

Figure A.1. Percentiles of actual experience conditional on potential experience



Note: Figure shows percentiles of actual against potential experience separately for men (Panel A) and women (Panel B). Actual experience is constructed from panel data for 1985–2014, while potential experience = age – years of education + 6. Solid line represents the 45-degree line, for which actual experience equals potential experience.

A.2 Additional Summary Statistics

Table A.1. Summary statistics over sample period

	2007			2014			Pooled 2007–2014		
	Overall	Men	Women	Overall	Men	Women	Overall	Men	Women
Share nonwhite	0.355	0.385	0.305	0.407	0.438	0.360	0.384	0.416	0.333
Share primary school	0.113	0.146	0.057	0.073	0.097	0.037	0.092	0.121	0.045
Share middle school	0.220	0.264	0.146	0.186	0.226	0.126	0.204	0.246	0.136
Share high school	0.475	0.460	0.500	0.554	0.548	0.563	0.520	0.508	0.540
Share college	0.192	0.130	0.296	0.186	0.129	0.274	0.183	0.125	0.278
Mean years of education	10.8	10.1	12.0	11.3	10.7	12.1	11.0	10.4	12.1
(std. dev.)	(3.5)	(3.5)	(3.2)	(3.1)	(3.2)	(2.8)	(3.3)	(3.4)	(3.0)
Mean age	33.5	33.2	34.0	34.5	34.3	34.9	33.9	33.6	34.2
(std. dev.)	(9.5)	(9.5)	(9.5)	(9.5)	(9.5)	(9.5)	(9.5)	(9.5)	(9.5)
Mean establishment size	3,477	1,877	6,179	3,627	2,540	5,284	3,135	1,860	5,216
(std. dev.)	(19,946)	(12,956)	(27,831)	(18,287)	(13,743)	(23,485)	(17,720)	(12,160)	(24,046)
Mean gender-establishment size	2,182	850	4,431	2,267	1,420	3,558	1,943	907	3,633
(std. dev.)	(12,989)	(3,518)	(20,607)	(11,217)	(5,959)	(16,144)	(11,101)	(3,924)	(17,165)
Mean establishment age	30.6	27.8	35.5	33.5	30.8	37.7	31.0	28.2	35.6
(std. dev.)	(22.5)	(21.2)	(23.8)	(23.6)	(22.4)	(24.7)	(23.1)	(21.9)	(24.4)
Mean months employed in year	9.9	9.7	10.1	10.1	10.0	10.3	9.9	9.8	10.1
(std. dev.)	(3.2)	(3.3)	(3.2)	(3.0)	(3.0)	(3.0)	(3.2)	(3.2)	(3.1)
Mean contractual work hours	41.3	42.4	39.5	41.3	42.3	39.8	41.5	42.5	39.8
(std. dev.)	(5.6)	(4.2)	(7.1)	(5.4)	(4.1)	(6.6)	(5.4)	(4.0)	(6.7)
Mean tenure (years)	4.3	3.7	5.2	4.4	4.1	4.8	4.1	3.7	4.8
(std. dev.)	(6.0)	(5.5)	(6.6)	(6.0)	(5.7)	(6.4)	(5.9)	(5.5)	(6.4)
Mean log real monthly earnings	7.108	7.156	7.028	7.367	7.430	7.272	7.237	7.291	7.150
(std. dev.)	(0.728)	(0.732)	(0.714)	(0.691)	(0.692)	(0.677)	(0.709)	(0.712)	(0.695)
Mean log real hourly wage	3.402	3.416	3.377	3.659	3.692	3.609	3.525	3.548	3.487
(std. dev.)	(0.780)	(0.769)	(0.797)	(0.742)	(0.730)	(0.757)	(0.758)	(0.747)	(0.774)
Number of worker-years	24,348,192	15,292,100	9,056,092	29,881,399	18,039,128	11,842,271	231,805,831	143,745,869	88,059,962
Number of unique workers	24,348,192	15,292,100	9,056,092	29,881,399	18,039,128	11,842,271	55,078,455	33,197,634	21,880,821
Number of unique establishments	184,168	135,346	48,822	191,504	138,171	53,333	222,695	153,081	69,614
Share female	0.372			0.396			0.380		
Mean log gender earnings gap	0.128			0.158			0.141		
Mean log gender wage gap	0.039			0.083			0.062		

Note: Table shows summary statistics for 2007, 2014, and the pooled sample 2007–2014 after sample selection and restriction to connected set.

Table A.2. Summary statistics before sample selection and restriction to connected set

	2007			2014			Pooled 2007–2014		
	Overall	Men	Women	Overall	Men	Women	Overall	Men	Women
Share nonwhite	0.339	0.368	0.294	0.388	0.416	0.349	0.365	0.394	0.321
Share primary school	0.113	0.149	0.056	0.072	0.099	0.035	0.090	0.122	0.044
Share middle school	0.232	0.278	0.160	0.189	0.233	0.129	0.209	0.255	0.143
Share high school	0.493	0.462	0.543	0.586	0.560	0.621	0.545	0.515	0.588
Share college	0.162	0.111	0.241	0.153	0.108	0.214	0.156	0.109	0.225
Mean years of education	10.7	10.0	11.7	11.1	10.5	11.9	10.9	10.3	11.8
(std. dev.)	(3.4)	(3.5)	(3.0)	(3.0)	(3.1)	(2.6)	(3.2)	(3.3)	(2.8)
Mean age	33.0	33.0	33.0	33.7	33.8	33.7	33.3	33.4	33.3
(std. dev.)	(9.5)	(9.5)	(9.4)	(9.5)	(9.6)	(9.5)	(9.5)	(9.5)	(9.5)
Mean establishment size	2,137	1,188	3,610	2,055	1,503	2,817	15,300	9,429	23,897
(std. dev.)	(15,571)	(10,183)	(21,317)	(13,672)	(10,416)	(17,144)	(110,234)	(76,567)	(145,728)
Mean gender-establishment size	1,347	542	2,596	1,291	849	1,900	9,522	4,622	16,697
(std. dev.)	(10,196)	(2,759)	(15,843)	(8,451)	(4,587)	(11,859)	(69,490)	(24,776)	(104,473)
Mean establishment age	23.5	22.0	25.8	23.8	22.6	25.6	23.5	22.2	25.4
(std. dev.)	(21.3)	(19.9)	(23.0)	(22.3)	(21.1)	(23.7)	(21.8)	(20.5)	(23.3)
Mean months employed in year	9.6	9.5	9.7	9.7	9.7	9.8	9.6	9.6	9.7
(std. dev.)	(3.4)	(3.4)	(3.4)	(3.3)	(3.3)	(3.3)	(3.3)	(3.3)	(3.4)
Mean contractual work hours	42.0	42.8	40.9	42.1	42.8	41.2	42.1	42.8	41.2
(std. dev.)	(5.0)	(3.8)	(6.2)	(4.7)	(3.7)	(5.7)	(4.7)	(3.7)	(5.8)
Mean tenure (years)	3.8	3.5	4.2	3.6	3.5	3.8	3.6	3.4	3.9
(std. dev.)	(5.3)	(5.0)	(5.7)	(5.2)	(5.0)	(5.4)	(5.2)	(5.0)	(5.5)
Mean log real monthly earnings	6.981	7.034	6.898	7.232	7.299	7.139	7.116	7.176	7.028
(std. dev.)	(0.683)	(0.692)	(0.661)	(0.639)	(0.648)	(0.614)	(0.663)	(0.673)	(0.638)
Mean log real hourly wage	3.254	3.285	3.205	3.501	3.549	3.434	3.385	3.425	3.326
(std. dev.)	(0.733)	(0.727)	(0.739)	(0.686)	(0.683)	(0.685)	(0.710)	(0.706)	(0.710)
Number of worker-years	38,401,131	23,359,048	15,042,083	50,798,080	29,427,684	21,370,396	364,776,727	216,756,022	148,020,705
Number of unique workers	38,401,131	23,359,048	15,042,083	50,798,080	29,427,684	21,370,396	77,297,426	44,401,043	32,896,383
Number of unique establishments	2,716,661	1,664,469	1,052,192	3,583,804	2,062,313	1,521,491	5,927,621	3,329,980	2,597,641
Share female	0.392			0.421			0.406		
Mean log gender earnings gap	0.137			0.161			0.148		
Mean log gender wage gap	0.080			0.115			0.099		

Note: Table shows summary statistics for 2007, 2014, and the pooled sample 2007–2014 before sample selection and restriction to connected set.

Table A.3. Comparison of summary statistics before and after sample selection and restriction to connected set

	Pooled 2007–2014, connected set			Pooled 2007–2014, all			Ratio: connected set to all		
	Overall	Men	Women	Overall	Men	Women	Overall	Men	Women
Share nonwhite	0.384	0.416	0.333	0.365	0.394	0.321	1.052	1.056	1.037
Share primary school	0.092	0.121	0.045	0.090	0.122	0.044	1.022	0.992	1.023
Share middle school	0.204	0.246	0.136	0.209	0.255	0.143	0.976	0.965	0.951
Share high school	0.520	0.508	0.540	0.545	0.515	0.588	0.954	0.986	0.918
Share college	0.183	0.125	0.278	0.156	0.109	0.225	1.173	1.147	1.236
Mean years of education	11.0	10.4	12.1	10.9	10.3	11.8	1.009	1.010	1.025
(std. dev.)	(3.3)	(3.4)	(3.0)	(3.2)	(3.3)	(2.8)	1.031	1.030	1.071
Mean age	33.9	33.6	34.2	33.3	33.4	33.3	1.018	1.006	1.027
(std. dev.)	(9.5)	(9.5)	(9.5)	(9.5)	(9.5)	(9.5)	1.000	1.000	1.000
Mean establishment size	25,078	14,878	41,728	15,300	9,429	23,897	1.639	1.578	1.746
(std. dev.)	(141,762)	(97,283)	(192,368)	(110,234)	(76,567)	(145,728)	1.286	1.271	1.320
Mean gender-establishment size	15,542	7,259	29,062	9,522	4,622	16,697	1.632	1.571	1.741
(std. dev.)	(88,805)	(31,395)	(137,318)	(69,490)	(24,776)	(104,473)	1.278	1.267	1.314
Mean establishment age	31.0	28.2	35.6	23.5	22.2	25.4	1.319	1.270	1.402
(std. dev.)	(23.1)	(21.9)	(24.4)	(21.8)	(20.5)	(23.3)	1.060	1.068	1.047
Mean months employed in year	9.9	9.8	10.1	9.6	9.6	9.7	1.031	1.021	1.041
(std. dev.)	(3.2)	(3.2)	(3.1)	(3.3)	(3.3)	(3.4)	0.970	0.970	0.912
Mean contractual work hours	41.5	42.5	39.8	42.1	42.8	41.2	0.986	0.993	0.966
(std. dev.)	(5.4)	(4.0)	(6.7)	(4.7)	(3.7)	(5.8)	1.149	1.081	1.155
Mean tenure (years)	4.1	3.7	4.8	3.6	3.4	3.9	1.139	1.088	1.231
(std. dev.)	(5.9)	(5.5)	(6.4)	(5.2)	(5.0)	(5.5)	1.135	1.100	1.164
Mean log real monthly earnings	7.237	7.291	7.150	7.116	7.176	7.028	1.017	1.016	1.017
(std. dev.)	(0.709)	(0.712)	(0.695)	(0.663)	(0.673)	(0.638)	1.069	1.058	1.089
Mean log real hourly wage	3.525	3.548	3.487	3.385	3.425	3.326	1.041	1.036	1.048
(std. dev.)	(0.758)	(0.747)	(0.774)	(0.710)	(0.706)	(0.710)	1.068	1.058	1.090
Number of worker-years	231,805,831	143,745,869	88,059,962	364,776,727	216,756,022	148,020,705	0.635	0.663	0.595
Number of unique workers	55,078,455	33,197,634	21,880,821	77,297,426	44,401,043	32,896,383	0.713	0.748	0.665
Number of unique establishments	222,695	153,081	69,614	5,927,621	3,329,980	2,597,641	0.038	0.046	0.027
Share female	0.380			0.406			0.936		
Mean log gender earnings gap	0.141			0.148			0.953		
Mean log gender wage gap	0.062			0.099			0.626		

Note: Table shows summary statistics for connected set, all observations, and their ratio for the pooled sample 2007–2014.

B Empirical Appendix

B.1 A Mincerian Approach to Measuring the Gender Pay Gap

As a starting point in our pursuit of understanding the sources of the gender pay gap, we run a series of classical Mincer regressions without controls for employer identity (Mincer, 1974; Heckman et al., 2006). The goal is twofold. First, we want to deliver a set of estimates that are directly comparable to the large existing literature that has studied gender gaps using household surveys or similar datasets. Second, we want to understand the part of the gender gap that is explained by worker and job characteristics associated with labor supply factors, which we think of as orthogonal to employer characteristics and other labor demand factors.

A classical Mincerian specification for income (i.e., either earnings or wages) of individual i in year t , denoted y_{it} , is simply

$$y_{it} = X_{it}\beta + \mathbf{1}[\text{gender}_i = M] \alpha^M + \mathbf{1}[\text{gender}_i = F] \alpha^F + \varepsilon_{it}, \quad (\text{B.1})$$

where X_{it} is a vector of observable worker and job characteristics discussed below; $\mathbf{1}[\text{gender}_i = M]$ and $\mathbf{1}[\text{gender}_i = F]$ are indicator functions that equal 1 if the gender of individual i is male or female, respectively, and 0 otherwise; α^M and α^F are gender-specific intercepts; and ε_{it} is a residual term. We estimate this equation via ordinary least squares (OLS) under the usual strict exogeneity assumption that $\mathbb{E}[\varepsilon_{it}|X_{it}, \text{gender}_i] = 0$. The main object of interest resulting from equation (B.1) is the (conditional) gender pay gap $\gamma \equiv \alpha^M - \alpha^F$, which captures the mean pay difference between female versus male workers who are otherwise observationally identical.

Table B.1 shows the (conditional) gender gap in four different specifications: the earnings gap without any controls in column 1; the earnings gap controlling for a linear term in years of education and a second-order polynomial in actual experience in column 2; the wage gap with the same controls in column 3; and the earnings gap with an additional set of dummies for education, actual experience, age, hours, nationality, municipality, industry, occupation, and tenure in column 4.⁴⁴

Our preferred specification is reported in column 4, with a conditional gender pay gap of around 12 log points. By including a rich set of observable worker and job characteristics as controls, this specification purges the raw data from various labor supply-related gender differences highlighted in the previous literature—see, for example, Goldin (2014) and Erosa et al. (2019). This specification flexibly controls for hours dummies with earnings as the dependent variable. If we restricted hours to enter linearly with coefficient one in this regression, then this would be identical to using the wage rate as the dependent variable. More generally, a complete set of hours FEs controls for nonconstant wage rates as a function of hours worked. Finally, it is worth noting that our preferred specification yields a high R^2 value of close to 70 percent, which suggests that we are controlling for a set of gender-pay-relevant characteristics with considerable explanatory power.⁴⁵

We conclude that a large gender pay gap remains within narrowly defined population subgroups defined by a rich set of covariates related to labor supply. With a large gap left unexplained by labor supply, we next turn to factors related to labor demand and ask, What is the role of gender-specific employer heterogeneity in explaining the gender pay gap?

⁴⁴Table B.2 repeats the same set of mincer reg regressions for the year 2007. Between 2007 and 2014, the raw gender pay gap (corresponding to column 1 of Table B.1) increased by a little less than 3 log points, while the conditional gap with our full set of controls (corresponding to column 4 of Table B.1) shows a decline of a little over 1 log point.

⁴⁵In univariate regressions and when gradually introducing controls in specification (B.1), we find that occupation and hours dummies account for a significant share of empirical pay variation in general and the gender gap in particular.

Table B.1. Estimates from Mincer regressions, 2014

	(1)	(2)	(3)	(4)
Gender gap	0.150	0.246	0.192	0.119
Income concept	Earnings	Earnings	Wage	Earnings
Education (linear)		✓	✓	
Education FEs				✓
Actual experience (polynomial)		✓	✓	
Actual experience FEs				✓
Age FEs				✓
Contractual work hours FEs				✓
Nationality FEs				✓
Municipality FEs				✓
Industry FEs				✓
Occupation FEs				✓
Tenure FEs				✓
Observations	31,830,960	31,830,960	31,830,960	31,830,960
R ²	0.012	0.372	0.389	0.698

Note: Table shows estimated (conditional) gender pay gap for various income concepts and controls based on estimating equation (B.1) on data from 2014. “Education (linear)” represents a linear term in years of education, “education FEs” represents dummies for nine education categories, “actual experience (polynomial)” represents a second-order polynomial in years of actual experience, “actual experience FEs” represents dummies for years of actual experience, “age FEs” represents dummies for years of age, “contractual work hours FEs” represents dummies for contractual work hours, “nationality FEs” represents dummies for worker nationality, “municipality FEs” represents dummies for approximately 5,500 establishment municipalities in Brazil, “industry FEs” represents dummies for 5-digit industry codes, “occupation FEs” represents dummies for 6-digit occupation codes, and “tenure FEs” represents dummies for years of tenure at current establishment.

B.2 Gender Segregation across Employers

Women make up 38 percent of Brazil’s formal sector employment over the period of 2007–2014. However, extending previous work on the U.S. (Hellerstein et al., 2011, Chapter 2 of Sorkin, 2015) and Portugal (Card et al., 2016), we find that women in Brazil are highly segregated across employers, even within industries and occupations. Figure B.1 shows a histogram of female employment shares in 2014. Around 28 percent of establishments have less than 10 percent women among their workforce. In contrast, if women were equally distributed across employers, we would see a single bar of height 10 in the 30–40 percent category.

Gender segregation is a robust empirical phenomenon over time, within sectors, and across different employer sizes. Figure B.2 repeats the same histogram of female employment shares shown in Figure B.1 separately for the years 2007 and 2014, which yields the same qualitative conclusion as before.

Figure B.3 splits the histogram shown in Figure B.1 into 9 sectors, which again yields the same qualitative conclusion as before.

Figure B.4 splits the histogram shown in Figure B.1 into 9 occupation groups, which again yields the same qualitative conclusion as before.


To show that the unequal distribution of women across employers is robust to different employer

Table B.2. Estimates from Mincerian regressions, 2007

	(1)	(2)	(3)	(4)
Gender gap	0.125	0.266	0.201	0.132
Income concept	Earnings	Earnings	Wage	Earnings
Education (linear)		✓	✓	
Education FEs				✓
Actual experience (polynomial)		✓	✓	
Actual experience FEs				✓
Age FEs				✓
Contractual work hours FEs				✓
Nationality FEs				✓
Municipality FEs				✓
Industry FEs				✓
Occupation FEs				✓
Tenure FEs				✓
Observations	25,208,660	25,208,660	25,208,660	25,208,632
R^2	0.007	0.367	0.385	0.682

Note: Table shows estimated (conditional) gender pay gap for various income concepts and controls based on estimating equation (B.1) on data from 2007. “Education (linear)” represents a linear term in years of education, “education FEs” represents dummies for nine education categories, “actual experience (polynomial)” represents a second-order polynomial in years of actual experience, “actual experience FEs” represents dummies for years of actual experience, “age FEs” represents dummies for years of age, “contractual work hours FEs” represents dummies for contractual work hours, “nationality FEs” represents dummies for worker nationality, “municipality FEs” represents dummies for approximately 5,500 establishment municipalities in Brazil, “industry FEs” represents dummies for 5-digit industry codes, “occupation FEs” represents dummies for 6-digit occupation codes, and “tenure FEs” represents dummies for years of tenure at current establishment.


Figure B.1. Histogram of female employment shares, 2007–2014



Note: Figure shows histogram of establishment-level female employment shares in ten bins (0–10 percent, 10–20 percent, . . . , 90–100 percent) using 2014 data.

size selection criteria, Figure B.5 shows percentiles of the female employment share distribution across

Figure B.2. Histograms of female employment shares, 2007 and 2014



Note: Figure shows histogram of establishment-level female employment shares in ten bins (0–10 percent, 10–20 percent, . . . , 90–100 percent) using data for 2007 (Panel A) and 2014 (Panel B).

establishments for various employment size cutoffs. This addresses the concern that at small establishments, employment of women may not be representative of the population because of the indivisibility of bodies in the data. The figure shows that empirically we are quite far away from proportionate-to-population representation of women, even at very large establishments with more than 10,000 employees. Even among establishments with at least 1,000 employees in the data, the female employment share varies vastly between 5 and 81 percent, going from the 5th percentile to the 95th percentile of the female employment share distribution.

To illustrate the nonuniform distribution of women relative to that of men across employers, Figure B.6 shows a histogram of male and female employment shares across firms ranked by their employment of the other gender. If the distribution of women (men) were a symmetric scaled version of that of men (women), one would see a strictly monotonic increasing density of density bins of one gender across employment ranks of the other gender. In contrast, the histograms in panels A–D show pronounced nonmonotonicities, suggesting that employers that have a relatively large mass of men do not necessarily also have a relatively large mass of women.

B.3 Gender Segregation Index

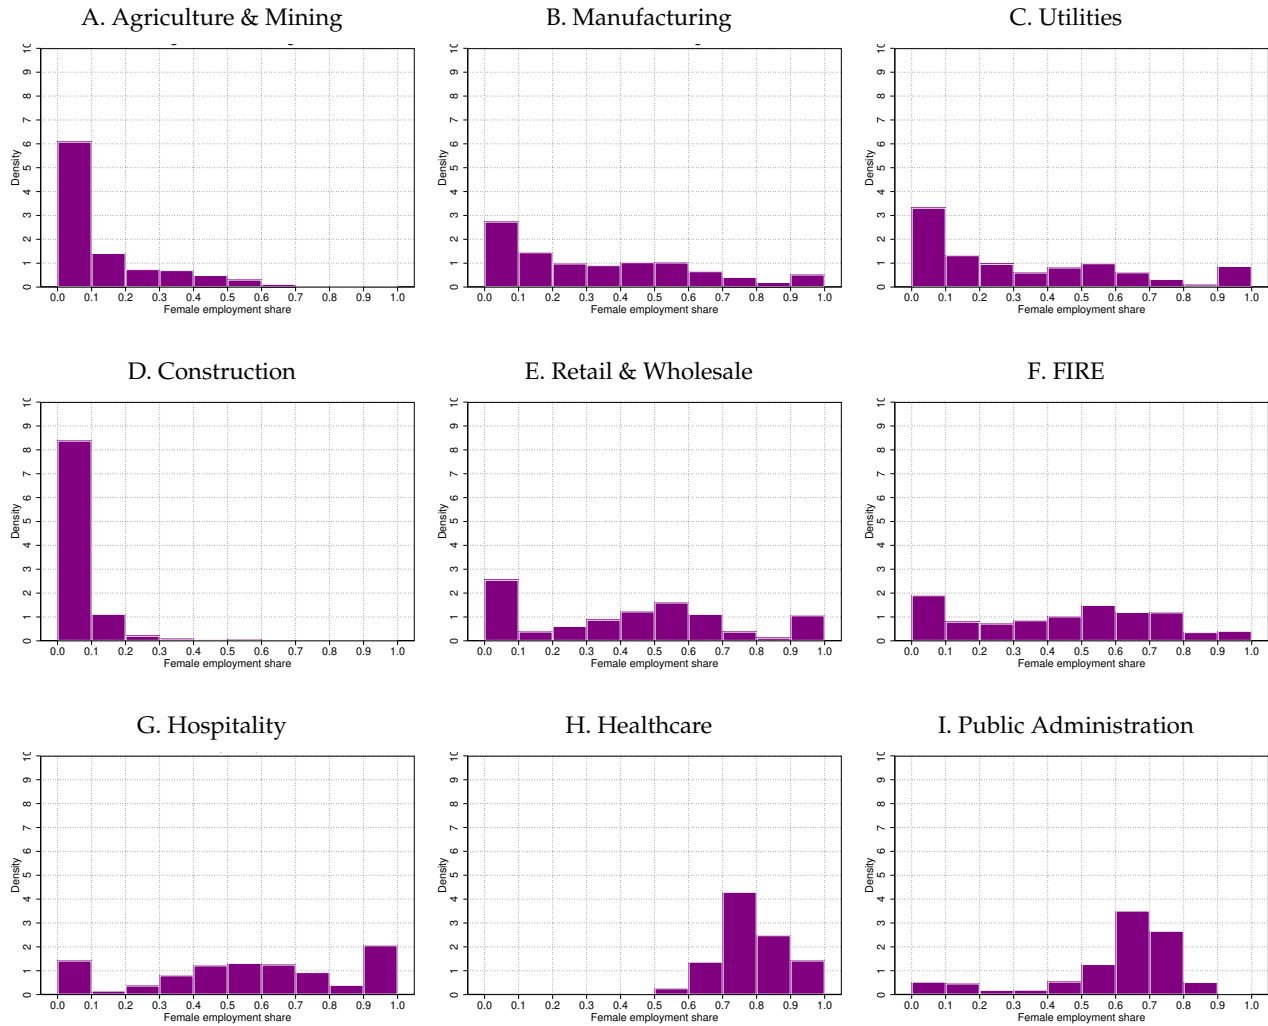
To quantify the extent to which women are nonuniformly distributed across employers, we define the following *employer segregation index*, S_t :

$$S_t = \frac{\sum_{i=1}^{N_m+N_f} \left((\text{firm-level female share})_{J(i,t)} - \text{population female share} \right)^2}{N_m \times (\text{population female share})^2 + N_f \times (1 - \text{population female share})^2}$$

where i indexes individual workers; N_m and N_f are the number of male and female workers, respectively; and $J(i, t)$ is a function that gives the index of the employer of individual i in year t . Note that the employer segregation index S_t lies between 0 and 1, with 0 meaning that each employer has a representative share of women and 1 meaning that all women work at employers where only women work (and, hence, similarly for men).

We find that the employer segregation index, S_t , takes on a value of 0.349 in 2007. To assess whether this is an economically meaningful deviation from uniform female shares, note that an index

Figure B.3. Histogram of female employment shares, by sector

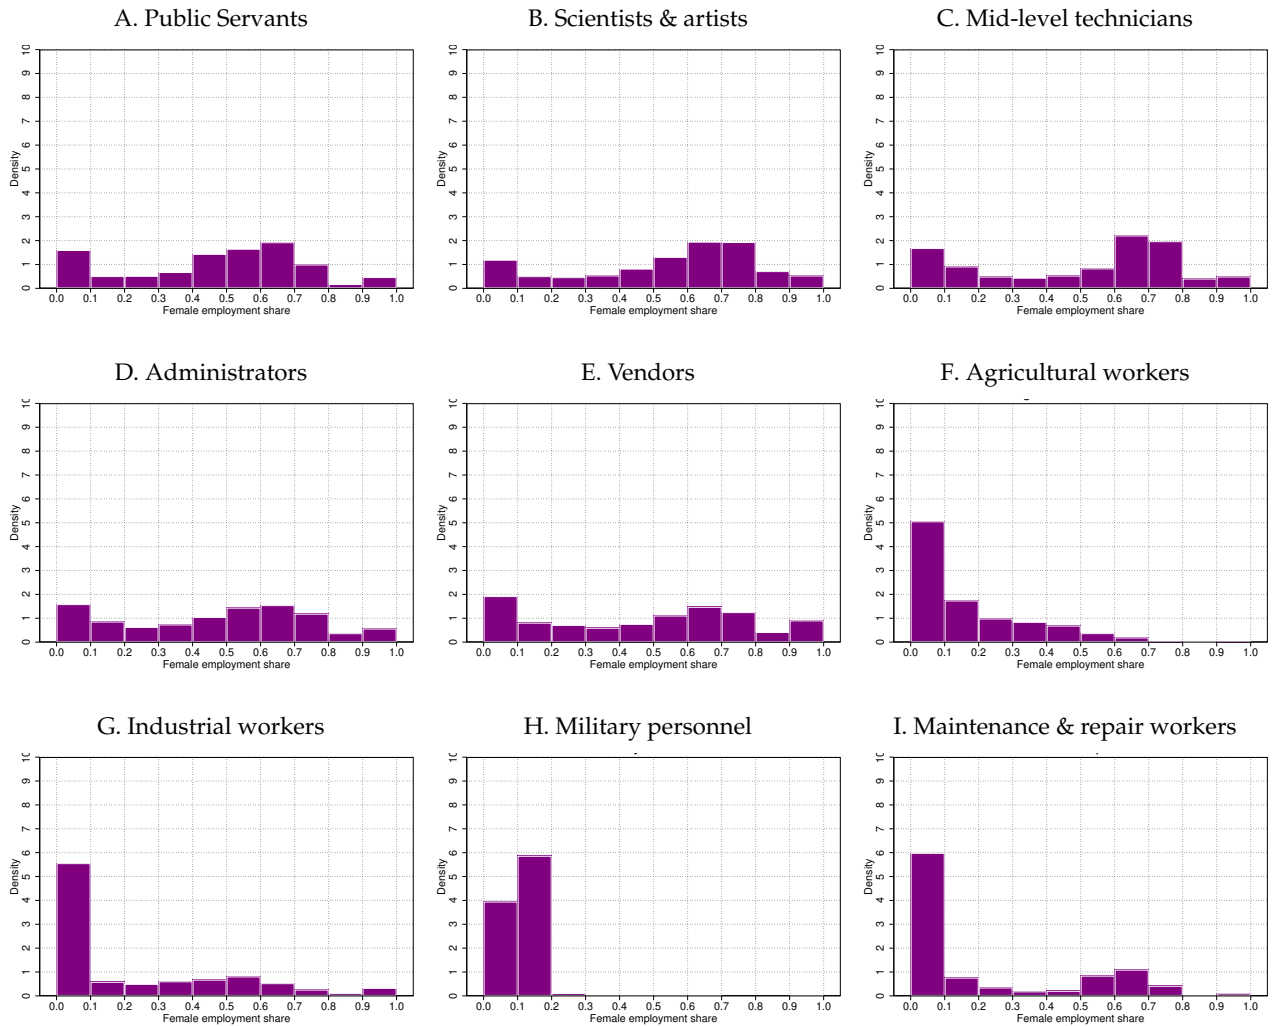


Note: Figure shows histograms of establishment-level female employment shares in ten bins (0–10 percent, 10–20 percent, . . . , 90–100 percent) within various one-digit sectors using 2014 data. FIRE stands for Finance, Insurance, and Real Estate.

value of 0.349 corresponds to an equivalent absolute value difference in employer gender shares of ± 0.288 around the population female share. In comparison, we find that in our data, the same index is significantly smaller when computed across industries (0.109), occupations (0.142), or states (0.002). Furthermore, the index value is relatively stable when we restrict attention to employers above minimum size thresholds of between 10 and 1,000 employees.

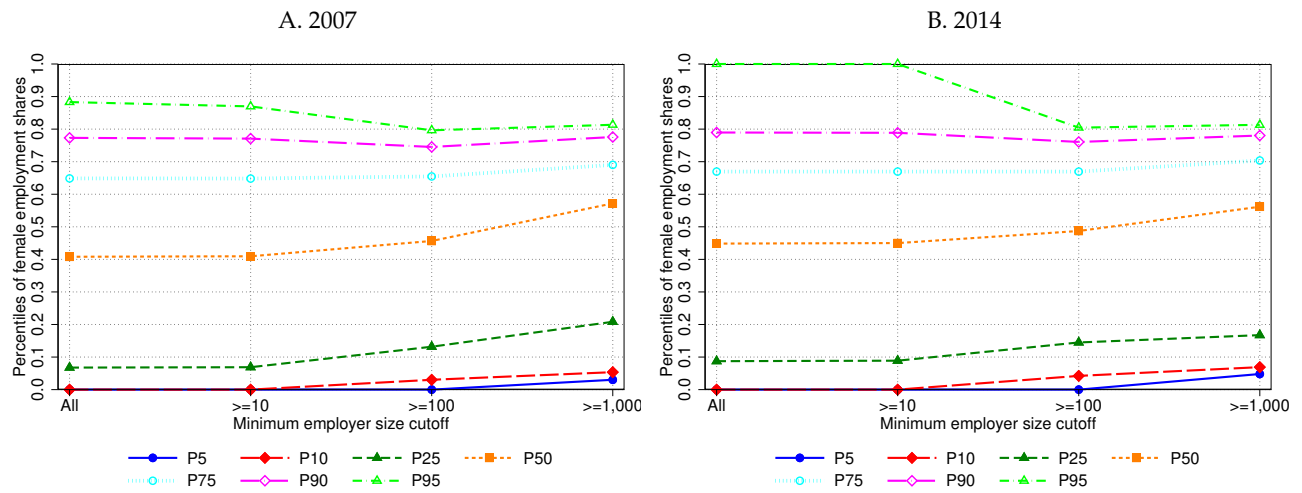
We conclude that there is a significant amount of gender segregation, with women distributed far from uniformly across employers.

Figure B.4. Histogram of female employment shares, by occupation



Note: Figure shows histograms of establishment-level female employment shares in ten bins (0–10 percent, 10–20 percent, . . . , 90–100 percent) within various one-digit occupation groups, using 2014 data.


Figure B.5. Percentiles of female employment share distribution




Note: Figure shows percentiles of female employment share distribution across establishments for the subpopulation of establishments falling above various minimum employer size cutoffs for 2007 (Panel A) and for 2014 (Panel B).

Figure B.6. Comparison of gender-specific employment distributions


A. Male distribution vs. female ranks, 2007




B. Male distribution vs. female ranks, 2014



C. Female distribution vs. male ranks, 2007




D. Female distribution vs. male ranks, 2014



Note: Figure shows histograms of male employment (Panels A and B) and female employment (Panels C and D) across the other gender’s employment rank distribution of establishments in ten bins (lowest ten percent of female employment, . . . , highest ten percent of female employment) using data from 2007 (Panels A and C) and 2014 (Panels B and D).


B.4 Further Details on AKM Estimation Results

Figure B.7. Predicted AKM contractual work hours FEs, by gender



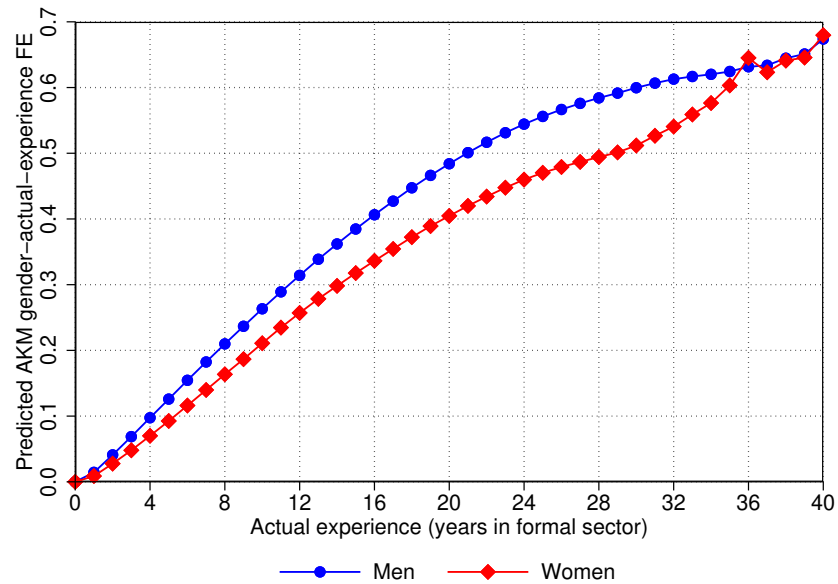
Note: Figure shows predicted AKM contractual work hours FEs separately for men and women based on estimating earnings equation (1). The omitted category is 1 hour, for which the FE value is normalized to 0.

Figure B.8. Predicted AKM occupation FEs, by gender



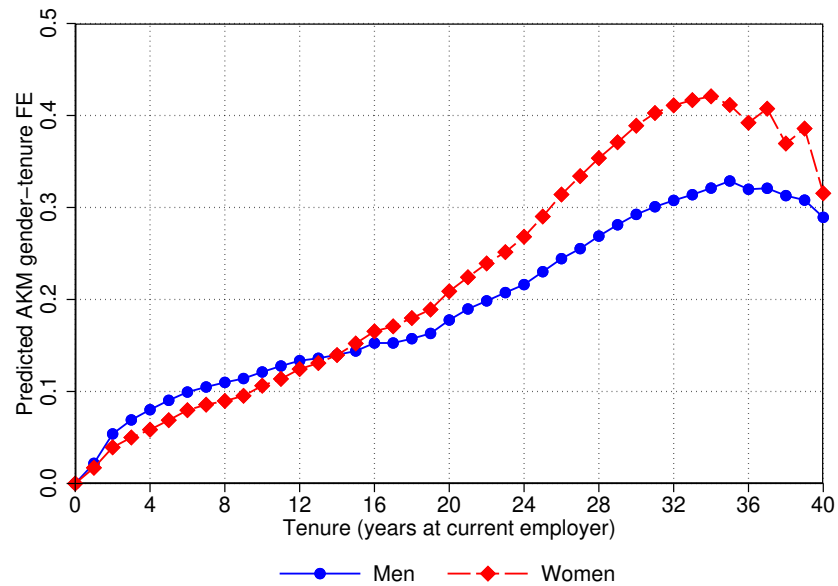
Note: Figure shows predicted AKM occupation FEs separately for men and women based on estimating earnings equation (1). Fixed effects of both genders are sorted by mean FEs of male FE quantiles.

Figure B.9. Predicted AKM actual-experience FEs, by gender




Note: Figure shows predicted AKM actual-experience FEs separately for men and women based on estimating earnings equation (1).

Figure B.10. Predicted AKM tenure FEs, by gender




Note: Figure shows predicted AKM tenure FEs separately for men and women based on estimating earnings equation (1).

Figure B.11. Predicted AKM education-year FEs, by gender



Note: Figure shows predicted AKM education-year FEs separately for men and women based on estimating earnings equation (1). Note that the declining pattern for both genders and all education categories is due to measuring earnings in multiples of the prevailing minimum wage, which increased over this period—see Engbom and Moser (2018) for details.

Figure B.12. Predicted AKM education-age FEs, by gender



Note: Figure shows predicted AKM education-age FEs separately for men and women based on estimating earnings equation (1). Age-pay profiles for all education groups are constrained to be constant from age 45 to age 49 and unconstrained otherwise.


B.5 Robustness of AKM Estimation Results

AKM demonstrated that estimates of employer FEs are still unbiased under standard assumptions. At the same time, high-dimensional two-way FE models of this type have been shown to suffer from limited-mobility bias in estimates of nonlinear forms, including variance and covariances of parameter estimates (Andrews et al., 2008, 2012; Bonhomme et al., 2020). Recent methodological advances have tackled this bias heads-on and proposed solutions (Kline et al., 2019; Bonhomme et al., 2019; Borovičková and Shimer, 2020). In this section, we provide evidence that in our setting and given our sample restrictions, limited-mobility bias is a lesser concern. Specifically, we show that the distributions of relevant parameter estimates are relatively stable over a range of employer size cutoffs between 10 employees – our baseline selection criterion – and levels as high as 1,000 employees.


The following figures show, separately by gender, the variance share of earnings due to estimated person FEs (Figure B.13), the variance share of earnings due to estimated employer FEs (Figure B.14), and the correlation between estimated person FEs and estimated employer FEs (Figure B.15). Comparing across employer size restrictions (colored bars) for a given subpopulation (horizontal axis), we note that the estimated moments vary somewhat between minimum employers size thresholds 1 and 10 but remain relatively stable thereafter. Thus, we conclude that our selected sample is sufficiently connected to not suffer from substantial limited-mobility bias. Consequently, one can reasonably trust our estimates of gender-specific employer FEs, which we use in our subsequent analysis.

Figure B.13. Variance share of earnings due to person FEs, by gender

A. Men




B. Women




Note: Figure shows variance share of earnings due to person FEs based on earnings equation (1) for various employer size restrictions in the data (colored bars) and subpopulations (horizontal axis) separately for men (Panel A) and women (Panel B).

Figure B.14. Variance share of earnings due to employer FEs, by gender

A. Men




B. Women




Note: Figure shows variance share of earnings due to employer FEs based on earnings equation (1) for various employer size restrictions in the data (colored bars) and subpopulations (horizontal axis) separately for men (Panel A) and women (Panel B).

Figure B.15. Correlation between person FEs and employer FEs, by gender

A. Men



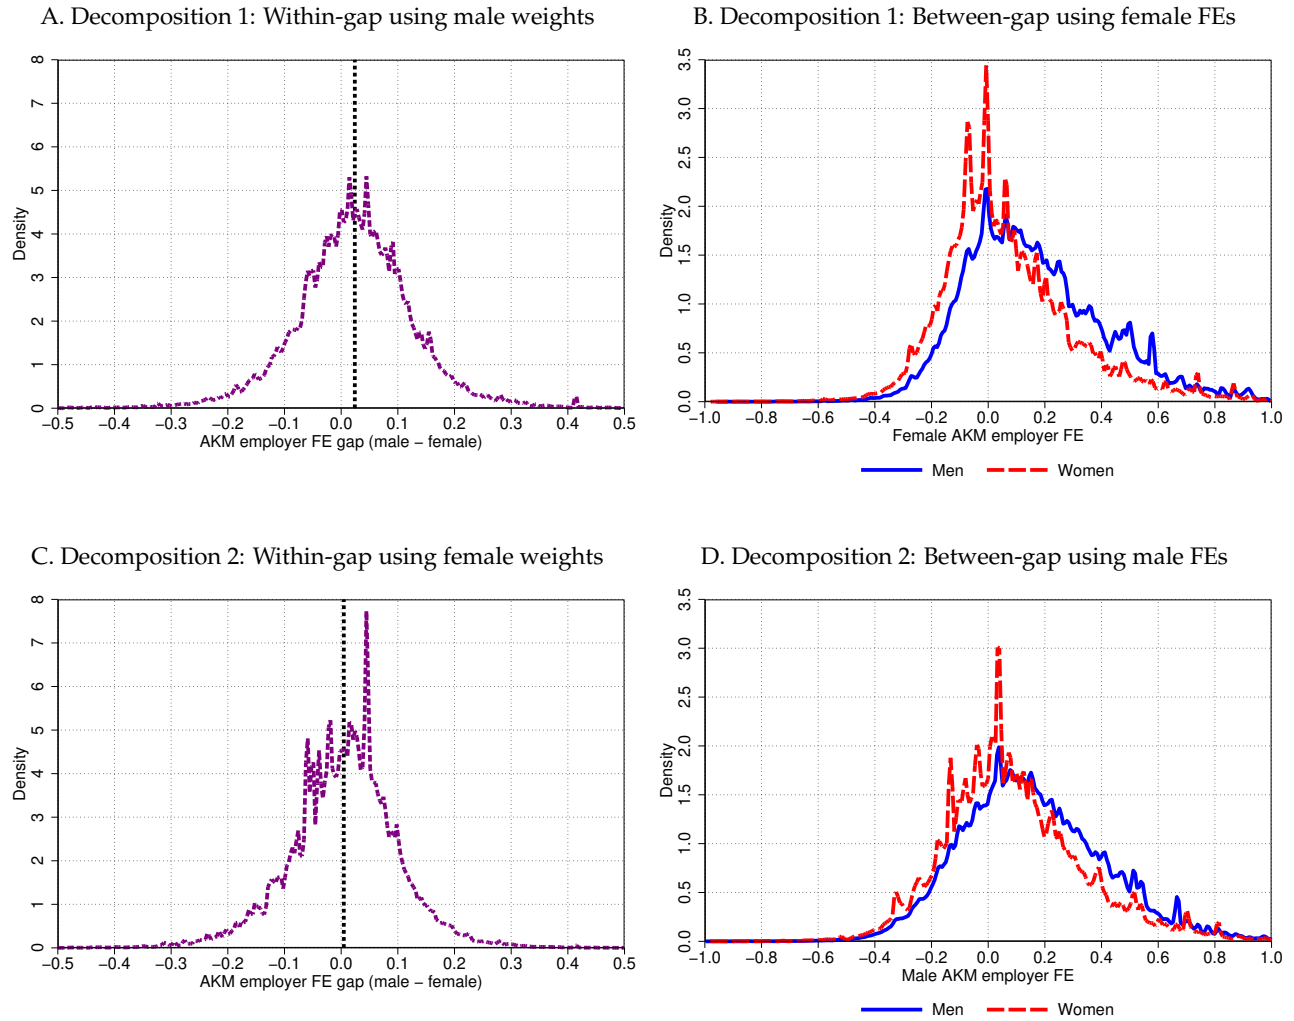
B. Women



Note: Figure shows correlation between person FEs and employer FEs based on earnings equation (1) for various employer size restrictions in the data (colored bars) and subpopulations (horizontal axis) separately for men (Panel A) and women (Panel B).

B.6 Further Details on Between vs. Within-Employer Pay Differences

Figure B.16. Components of Oaxaca-Blinder decompositions



Note: Figure shows pay distributions underlying the Oaxaca-Blinder decompositions—specifically, the within-gap using male weights (Panel A) and using female weights (Panel C) and the between-gap using female FEs (Panel B) and using male FEs (Panel D). Decomposition 1 (Panels A and B) and decomposition 2 (Panels C and D) correspond to equations (2) and (3) of the main text, respectively. Dashed vertical line shows mean of the distribution.

B.7 Life-Cycle Profiles by Gender and Parent Status

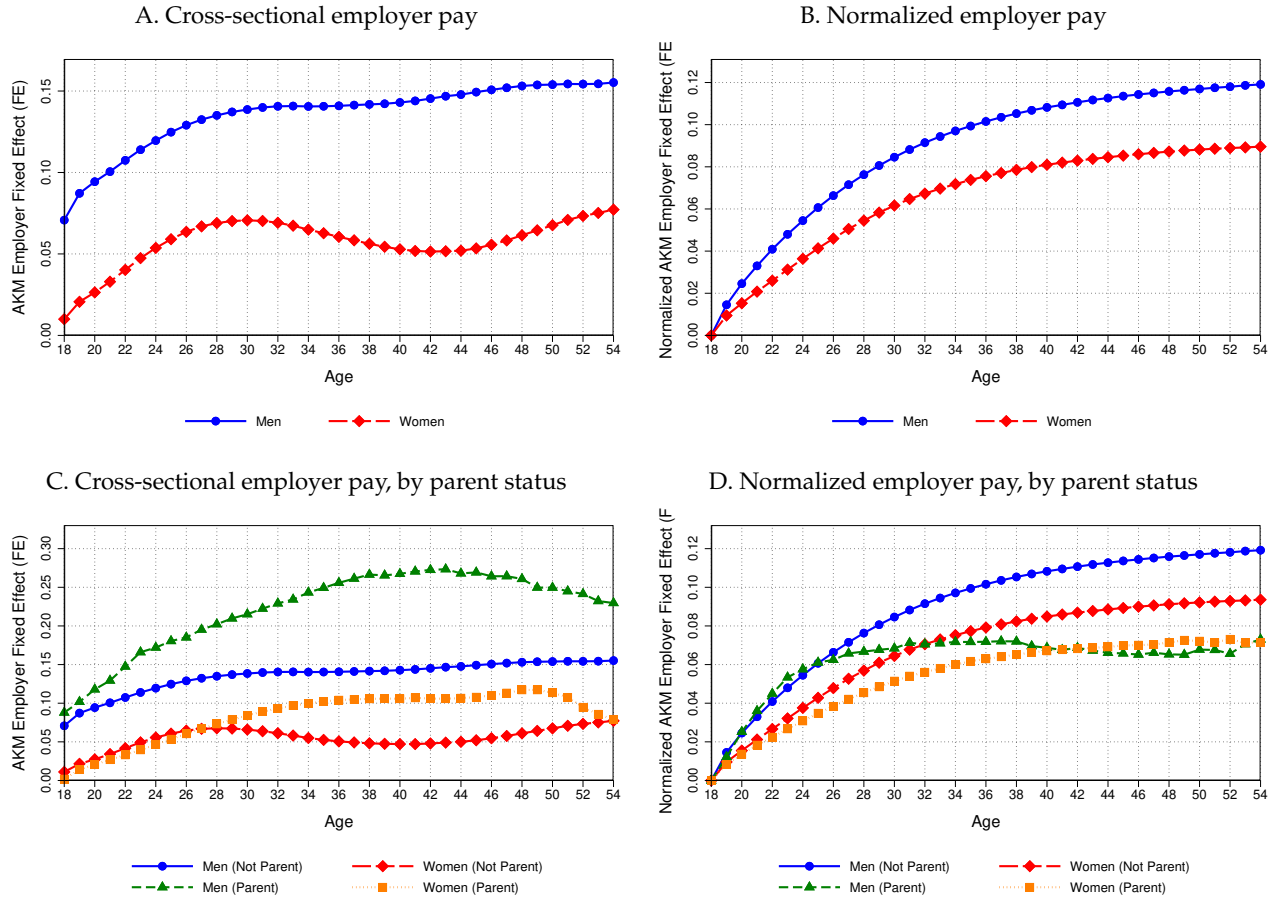
In this section, we are interested in life-cycle patterns in employer heterogeneity and how they differ by gender and parental status.⁴⁶ We compute two types of life-cycle statistics. The first set of statistics comprises raw, cross-sectional binned means. The second set of statistics comprises binned means of differenced variables, which we normalize to 0 at age 18.

Figure B.17 shows estimated gender-specific employer FEs by gender and parent status. A few things are worth noting. First, employer pay for women is less than that for men over the entire life-cycle. Second, cross-sectional life cycles (Panels A and C) can be quite different from the normalized life cycles (Panels B and D), plausibly owing to cohort effects and other dimensions of permanent individual heterogeneity that is differenced out in the normalized statistics. Third, both men and women see marked growth in employer FEs over their life cycle, although men significantly more so than women (Panel B). Fourth, parent men look more similar to women in general, and to women with children in particular, compared with nonparent men, although nonparent women still differ from nonparent men (Panel D).

Altogether, these life-cycle patterns suggest that childbirth could play some role in explaining parts of the gender pay gap, consistent with findings from similar studies in other contexts, such as Coudin et al. (2018) for France.

⁴⁶We classify individuals as “parent” if they ever went on registered parental leave from their employer during the sample period 2007–2014 and as “not parent” if they did not.

Figure B.17. Life-cycle mean gender-specific employer FEs, by gender and parent status



Note: Figure shows cross-sectional (Panels A and C) and normalized (Panels B and D) employer pay across age separately by gender (Panels A and B) and separately by gender and ever-parent status (Panels C and D). Cross-sectional estimates represent binned means. Normalized estimates are binned means of differenced variable, normalized to 0 at age 18.

B.8 Event Study Analysis around Parental Leaves by Gender

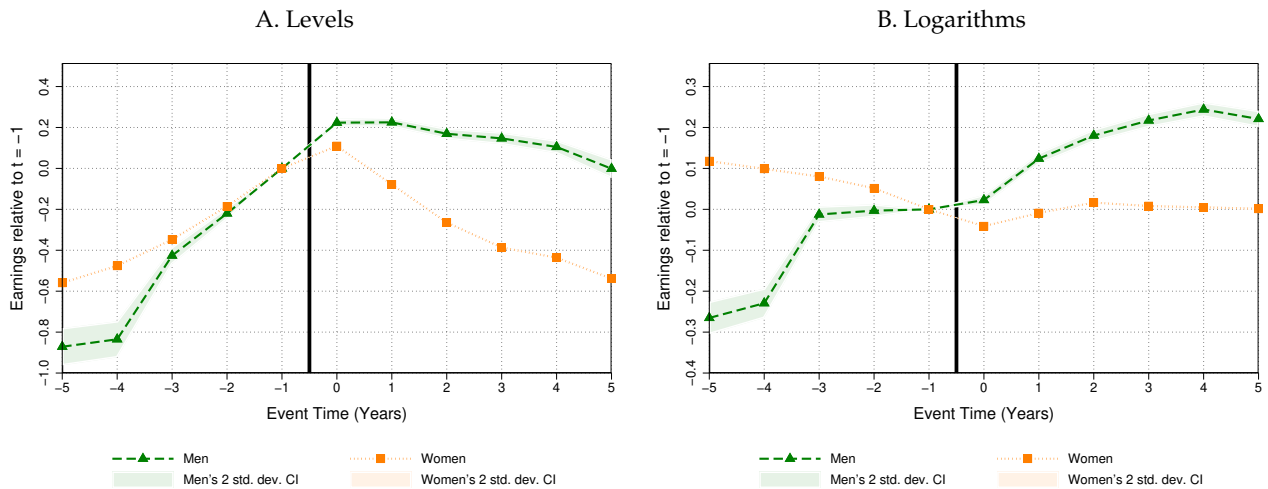
Following [Kleven et al. \(2016\)](#), we estimate the following event-study regression for individual i of gender g in year s and at event time t :

$$y_{ist} = \sum_{t' \neq 1} \alpha_{t'}^g \mathbf{1}[t' = t] + \sum_a \beta_a^g \mathbf{1}[a = \text{age}_{is}] + \sum_{s'} \gamma_{s'}^g \mathbf{1}[s' = s] + v_{ist}^g, \quad (\text{B.2})$$

where y_{ist} is the outcome variable of interest, $\alpha_{t'}^g$ denotes a set of gender-specific event time controls, β_a^g denotes a set of gender-specific age controls, $\gamma_{s'}^g$ denotes a set of gender-specific time controls, and v_{ist}^g is an error term. As dependent variables, we will use the level of earnings (filling in zero earnings for missing observations) or, alternatively, log earnings (dropping missing observations). Our focus will be on estimates of the coefficients $\alpha_{t'}^g$, based on equation (B.2), for men and women in an 11-year window around the birth of individuals' first child.

Figure B.18 plots the resulting event study graph, including gender-specific point estimates and confidence intervals. Panel A of the figure shows the event study for earnings in levels. Men and women are on comparable earnings paths leading up to their first child's birth, marked by the vertical black solid line. After their first child, women's earnings markedly decline, both in absolute value and compared with men's earnings, which remain relatively more stable. Panel B shows the event study for earnings in logarithms. Women's earnings show a declining pretrend in the five years leading up to their first child, both in absolute terms and compared to men, whose earnings increase over the same preperiod. After the birth of their child, women's earnings take a one-year dip and then remain relatively constant over the next five years. In contrast, men's earnings grow over the five years following their first child.

Figure B.18. Event-study plot of earnings relative to the year before the first child




Note: Figure shows event study tracking earnings levels (Panel A) and log earnings (Panel B) in a time window of 10 years around first childbirth. The vertical solid black line separates years before and after first childbirth.

Taken at face value, these results suggest that having children has an effect on the earnings and participation of women relative to men. However, it seems that women's earnings losses around having children are not systematically related to changes in the employer component of earnings. Figure B.19 illustrates this point by plotting an analogous event study with estimated gender-specific employer FEs from equation (1) as the dependent variable. Men and women follow a similar trend before the birth of a child. In the first two years after the birth of a child, women's gender-specific

employer FE falls behind that for men, but the pattern reverses during years 3 through 5. At any time in the event study, the gender gap in gender-specific employer FEs is less than 1 log point.

Altogether, this suggests that firm pay heterogeneity is *not* not a very important factor behind women’s pay penalty for having children.

Figure B.19. Event-study plot of gender-specific employer FEs rel. to the year before the birth of the first child



Note: Figure shows event study tracking gender-specific employer FEs based on estimating earnings equation (1) on the 2007–2014 sample in a time window of 10 years around the birth of a first child. The vertical solid black line separates years before and after the birth.

B.9 Details on Construction and Comparison of Employer Rank Measures

In this section, we define and implement alternative employer rank measures, which we then compare with the PageRank used in the main section of the paper. All three employer rank measures are consistent with a large class of on-the-job search models, including the structural framework that we will develop later on. For notational convenience, in this section, j by $r^{g,Page}(j)$ will denote the PageRank of an employer.

Poaching rank. According to the poaching rank (Moscarini and Postel-Vinay, 2008; Bagger and Lentz, 2019), higher-ranked employers hire relatively more workers from employment than from unemployment. Formally, the *poaching index* is defined as the share of all new hires that are due to poached workers from other employers:

$$s^{g,poach}(j) = \frac{n^g(.,j)}{n^g(0,j) + n^g(.,j)}, \quad (\text{B.3})$$

where $n^g(.,j)$ is the number of gender-specific hires that employer j makes from employment at other establishments and $n^g(0,j)$ is the number of gender-specific hires that employer j makes from unemployment. Intuitively, if the underlying employer rank of an establishment is higher, then it poaches more workers from its competitors, so the poaching index is increasing in the underlying employer rank. Finally, we construct the *poaching rank* of an employer as its rank among the set of poaching indices constructed as in equation (B.3), with the lowest rank normalized to 0 and the highest rank normalized to 100:

$$r^{g,poach}(j) = 100 \frac{\sum_{j' \in \mathcal{J}^g} \mathbf{1} [s^{g,poach}(j') \leq s^{g,poach}(j)]}{N^g}. \quad (\text{B.4})$$

Net poaching rank. According to the net poaching rank (Haltiwanger et al., 2018; Moscarini and Postel-Vinay, 2018), higher-ranked employers hire relatively more workers from other competitors and lose relatively fewer workers to other competitors. Formally, the *net poaching index* is defined as the net growth rate of an establishment's employment due to job-to-job transitions into and out of it:

$$s^{g,net}(j) = \frac{n^g(.,j) - n^g(j,.)}{E^g(j)}, \quad (\text{B.5})$$

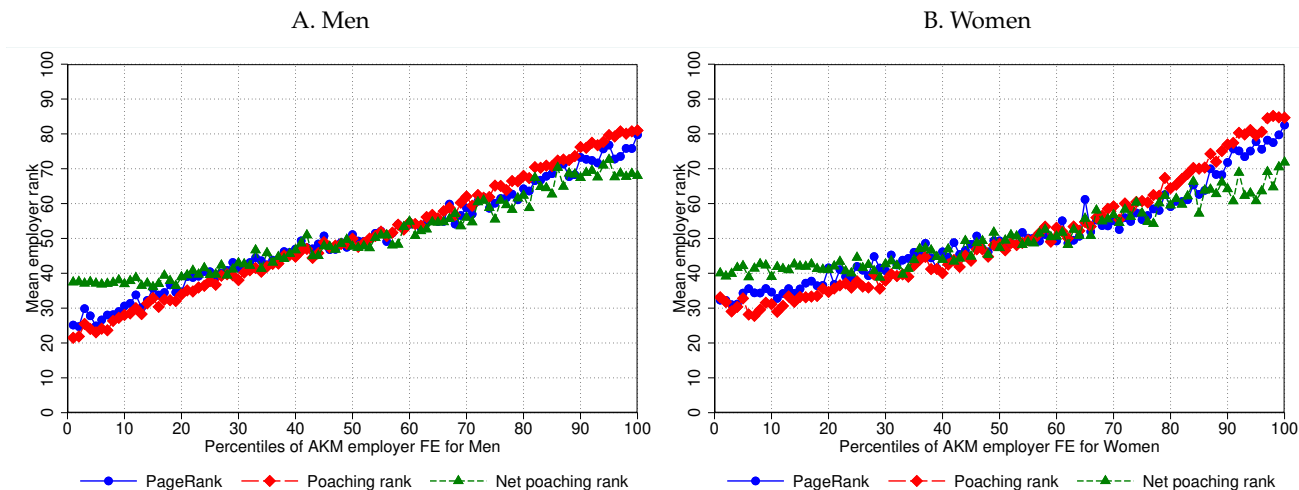
where $n^g(.,j)$ is the number of hires that employer j makes from employment at other establishments, $n^g(j,.)$ is the number of workers that employer j loses to other establishments through job-to-job transitions, and $E^g(j)$ is the gender-specific size of the workforce of employer j . Intuitively, if the underlying employer rank of an establishment is higher, then it poaches more workers from its competitors and retains more of its own workers, so the net poaching index is increasing in the underlying employer rank. Finally, we construct the *net poaching rank* of an employer as its rank among the set of net poaching indices constructed as in equation (B.5), with the lowest rank normalized to 0 and the highest rank normalized to 100:

$$r^{g,net}(j) = 100 \frac{\sum_{j' \in \mathcal{J}^g} \mathbf{1} [s^{g,net}(j') \leq s^{g,net}(j)]}{N^g}. \quad (\text{B.6})$$

Comparison of alternative employer rank measures. To compare the PageRank from the main text with the poaching rank from equation (B.4) and the net poaching rank from equation (B.6), Figure B.20 shows the relationship of the means of the three employer rank measures with the estimated

employer pay FEs by gender. There is a strong positive correlation between all three indices across genders. Although they are not perfectly correlated, particularly in the tails, the overall shape and slope across employer FE ranks is remarkably similar.

Figure B.20. Comparison of employer rank measures



Note: Figure compares various gender-specific employer rank measures for men (Panel A) and for women (Panel B). Employer rank measures include the PageRank, the poaching rank, and the net poaching rank.

While all three employer rank measures are strongly related on average, there are also some important discrepancies between them. Table B.3 shows rank correlations between the three employer rank measures and also their pay rank by gender. The correlation between PageRank and poaching rank is 0.549 for men and 0.552 for women. That between between PageRank and net poaching rank is 0.226 for men and 0.236 for women. All three measures are positively related to pay rank, with correlations between 0.256 and 0.469 for men and between 0.242 and 0.448 for women. Note that the rank correlation between the net poaching rank and other employment ranks, as well as pay rank, is relatively weak suggesting a larger role for nonpay employer characteristics in explaining the data.

Table B.3. Rank correlations of various employer rank measures, by gender

	Men				Women			
	PageRank	Poaching rank	Net poaching rank	Pay rank	PageRank	Poaching rank	Net poaching rank	Pay rank
PageRank	1.000				1.000			
Poaching rank	0.549	1.000			0.552	1.000		
Net poaching rank	0.226	0.425	1.000		0.236	0.430	1.000	
Pay rank	0.400	0.469	0.256	1.000	0.359	0.448	0.242	1.000

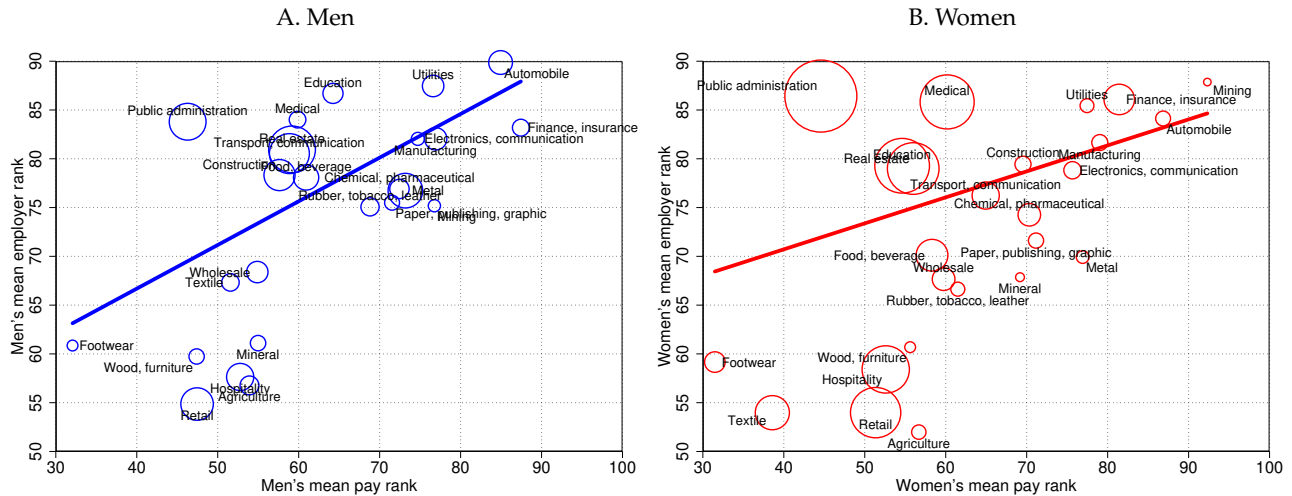
Note: Table shows correlations between various gender-specific employer rank measures separately for men and for women.

B.10 Further Details on Fact 2

Employer ranks versus pay by industry. We find interesting heterogeneity for both genders in pay and employer ranks across industries. Figure B.21 shows the mean pay ranks and mean employer ranks across 25 industries, with circle sizes representing employment shares and the solid line showing the weighted linear best fit. Comparing panel A and panel B, the aggregate positive correlation

between employer ranks and pay ranks is again evident, as is the lower gradient for women compared to men. There are many similarities between industry-level mean pay ranks and mean employer ranks across genders. For example, Footwear is the lowest-paying of all sectors, while Retail and Agriculture are the lowest-ranked sectors. Sectors that are high paying and attractive for both genders include Finance and Insurance, Utilities, and the Automobile sector. However, there are also interesting differences across genders. For example, Agriculture, Metal, and the Rubber, Tobacco, and Leather sector are relatively preferred by men, while the Medical sector and Public Administration are relatively preferred by women.

Figure B.21. Employer ranks versus pay across industries, by gender



Note: Figure shows gender-specific mean employer ranks using PageRank estimates, based on equation (4), against mean employer pay ranks using gender-specific pay FE estimates, based on earnings equation (1), across industries for men (Panel A) and for women (Panel B). Circle size is proportional to employment share. Solid line is weighted linear best fit.

Regression analysis of employer rank-pay relationship. Table B.4 shows the results of regressions of employer rank on employer pay rank with various controls. Columns (1) and (4) repeat the raw employer rank-pay rank relationship from above, which shows a gradient of 0.401 for men and 0.323 for women. In columns (2) and (5), five-digit industry FEs are added as controls, which reduces the gradient for both genders and somewhat more so for men, resulting in gradients of 0.364 for men and 0.316 for women. Finally, columns (3) and (6) add municipality FEs as controls. In this richest specification, the gradients are reduced further and by approximately the same absolutely amount for both genders, resulting in a gradient of 0.314 for men and 0.255 for women. We conclude that there remains a positive correlation between employer rank and pay rank that is steeper for men than for women, even within narrowly defined industries and geographic units.

Changes in pay by type of employer rank transition. In line with a job-ladder view of the world, Table B.5 shows the conditional changes in earnings and in establishment FEs upon making an employment-to-employment (henceforth “E-to-E”) transition between two consecutive years. We see that both men and women on average see an increase in earnings (employer FEs) upon making an E-to-E transition, and disproportionately so when moving up the employer rank distribution. However, the share of workers who see an increase in their earnings (employer FEs) upon transitioning is far below 1—namely, around 0.591 (0.630) for men and 0.586 (0.619) for women. Furthermore, the absolute change

Table B.4. Employer rank-pay rank gradient, various controls

	Men			Women		
	(1)	(2)	(3)	(4)	(5)	(6)
Employer pay rank	0.401*** (0.001)	0.364*** (0.001)	0.314*** (0.002)	0.323*** (0.002)	0.316*** (0.002)	0.255*** (0.002)
Industry FEs		✓	✓		✓	✓
Municipality FEs			✓			✓
Observations	143,745,8690	143,745,8690	143,745,8690	88,059,962	88,059,962	88,059,962
R ²	0.191	0.396	0.461	0.124	0.438	0.529

Note: Table shows estimated coefficient from regressing employer ranks using PageRank estimates, based on equation (4), against employer pay ranks using gender-specific pay FE estimates, based on earnings equation (1), separately by gender. ***, **, and * denote significance at the 1 percent, 5 percent, and 10 percent levels.

and the share of transitions with positive changes in earnings (employer FEs) are higher for men than for women. From this, we conclude that pay is positively but imperfectly correlated with employer ranks and that for men, pay is more closely aligned with employer ranks than it is for women.

Table B.5. Changes in earnings and employer FEs upon E-to-E transitions, by gender


	Men			Women		
	Overall	To lower rank	To higher rank	Overall	To lower rank	To higher rank
$\mathbb{E} [\Delta y_{ijt}]$	0.031	-0.018	0.072	0.033	-0.008	0.068
$\mathbb{E} [\mathbf{1} [\Delta y_{ijt} > 0]]$	0.549	0.500	0.591	0.552	0.512	0.586
$\mathbb{E} [\Delta \psi_j]$	0.021	-0.033	0.066	0.022	-0.024	0.060
$\mathbb{E} [\mathbf{1} [\Delta \psi_j > 0]]$	0.540	0.435	0.630	0.541	0.449	0.619
Share	0.206	0.095	0.112	0.128	0.058	0.070
Share of E-to-E transitions	1.000	0.459	0.541	1.000	0.456	0.544

Note: Table shows mean change in earnings y_{ijt} and estimated gender-specific employer FE in earnings based on earnings equation (1) of worker i employed at establishment j in year t between consecutive years.

B.11 Further Details on Fact 3

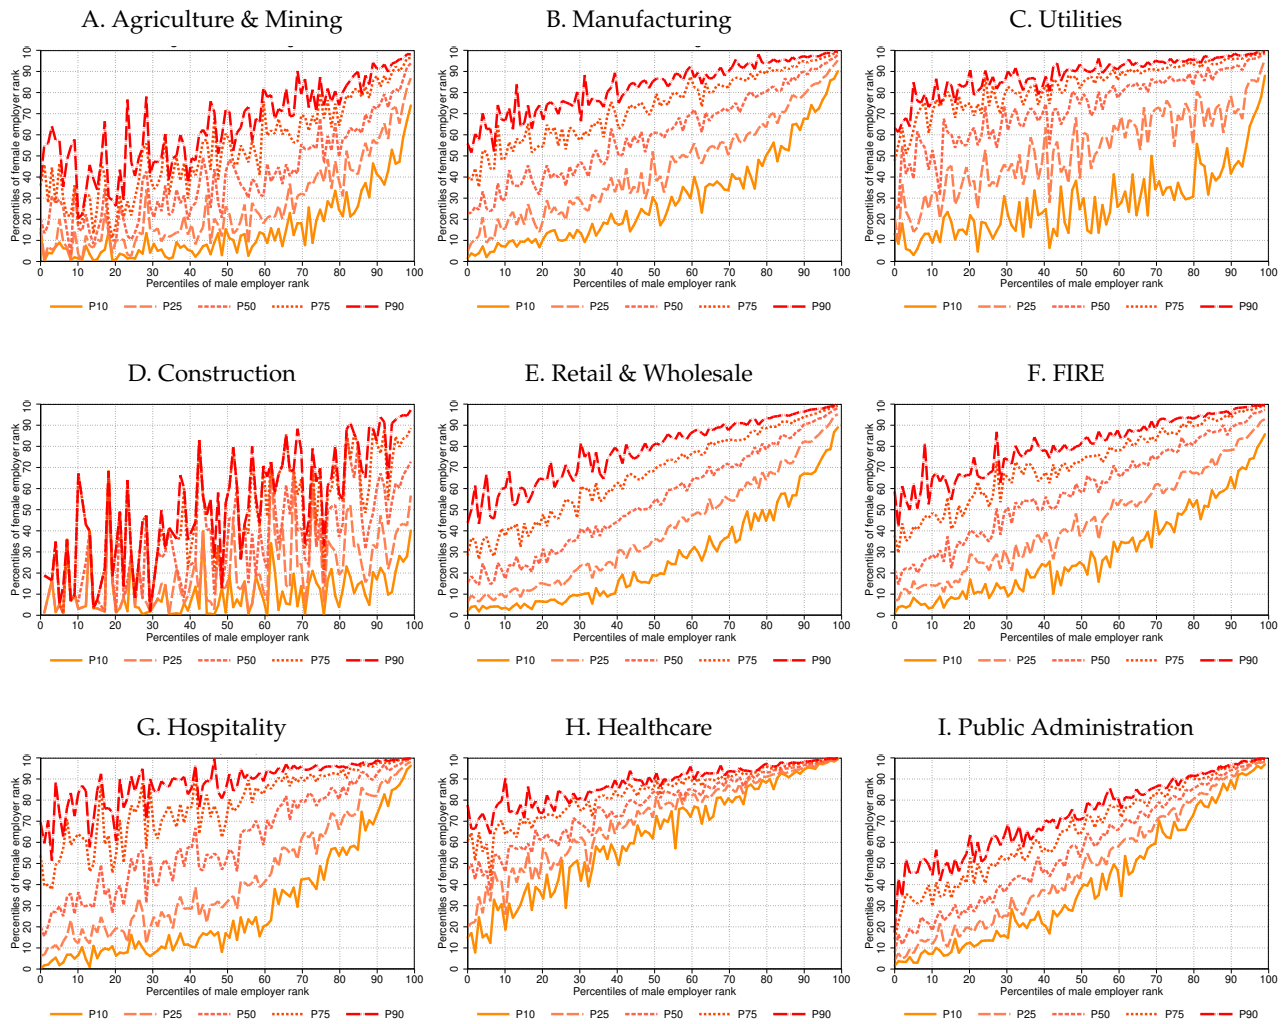
Cross-gender comparisons of pay and employer ranks.

Figure B.22. Female vs. male pay, by industry



Note: Figure shows various percentiles of women’s employer pay FE distribution, based on earnings equation (1), across quantiles of the male employer FE distribution across one-digit sectors. FIRE stands for Finance, Insurance, and Real Estate.

Figure B.23. Female vs. male employer ranks, by industry



Note: Figure shows various percentiles of women’s employer rank distribution using PageRank estimates, based on equation (4), across quantiles of the male employer rank distribution across one-digit sectors. FIRE stands for Finance, Insurance, and Real Estate.

C Model Appendix

C.1 Definition of a Stationary Equilibrium

We are now ready to define a *stationary equilibrium* for this economy.

Definition. A stationary search equilibrium is a set of worker value functions $\{S_{a,g}, W_{a,g}\}_{a,g}$ and policy functions $\{\phi_{a,g}\}_{a,g}$; firm value function Π and policy functions $\{w_{a,g}, \pi_{a,g}, v_{a,g}\}_{a,g}$; flow-utility offer distributions $\{F_{a,g}(x)\}_{a,g}$; measures of unemployed workers $\{u_{a,g}\}_{a,g}$, aggregate job searchers $\{U_{a,g}\}_{a,g}$, aggregate vacancies $\{V_{a,g}\}_{a,g}$, and labor market tightnesses $\{\theta_{a,g}\}_{a,g}$; job offer arrival rates $\{\lambda_{a,g}^u, \lambda_{a,g}^e, \lambda_{a,g}^G\}_{a,g}$; and firm sizes $\{l_{a,g}\}_{a,g}$ such that for all (a, g)

- given $F_{a,g}(x)$ and $\{\lambda_{a,g}^u, \lambda_{a,g}^e, \lambda_{a,g}^G\}$, the value functions $S_{a,g}$ and $W_{a,g}$ satisfy equations (5) and (6);
- unemployed workers' job acceptance policy follows a threshold rule $\phi_{a,g}$ given by equation (7), and employed workers with flow utility x accept any job x' such that $x' > x$;
- given $l_{a,g}$, firms' value function Π is given by equation (9);
- firms' policy functions $\{w_{a,g}, \pi_{a,g}, v_{a,g}\}$ solve the problem in equation (9);
- measures of unemployed workers are given by equation (8), aggregate job searchers $U_{a,g}$ are given by equation (10), aggregate vacancies $V_{a,g}$ are given by equation (11), and labor market tightness $\theta_{a,g}$ is given by equation (12);
- given $\theta_{a,g}$, the job offer arrival rates $\{\lambda_{a,g}^u, \lambda_{a,g}^e, \lambda_{a,g}^G\}$ satisfy equation (13);
- given $F_{a,g}(x)$, $\{\lambda_{a,g}^u, \lambda_{a,g}^e\}_{a,g}$, $\lambda_{a,g}^G$, and $V_{a,g}$, firm sizes satisfy equation (14);
- the offer distribution satisfies $F_{a,g}(x) = \int_j v_{a,g}(j) \mathbf{1}[x_{a,g}(j) \leq x] d\Gamma(j) / V_{a,g}$.

C.2 Additional Proofs

C.2.1 Proof of Lemma 1 (Optimal Amenities)

Proof. Based on the insight that workers care only about the flow utility of a job, we can rewrite the problem of a firm in equation (9) as one of choosing in each market a flow utility $x = w + \pi$ and vacancies v that solve the following problem:

$$\max_{x,v} \left\{ \left[pa - c_{a,g}^x(x) - z_{a,g} \right] l_{a,g}(x,v) - c_{a,g}^v(v) \right\}, \quad \forall (a,g),$$

where $c_{a,g}^x(x)$ is the solution to the following cost-minimization subproblem in each market:

$$c_{a,g}^x(x) = \min_{w,\pi} \left\{ w + c_{a,g}^\pi(\pi) \right\} \quad \text{s.t.} \quad w + \pi = x. \quad (\text{C.1})$$

Once written in this way, it is evident that an interior solution to the firm's cost-minimization problem in equation (C.1) is characterized by the following optimality conditions:

$$c_{a,g}^{\pi,0} \times \frac{\partial \tilde{c}_a^\pi(\pi^*)}{\partial \pi} = 1 \quad (\text{C.2})$$

$$w^* = x - \pi^*, \quad \forall (a,g).$$

Equation (C.2) uniquely pins down a firm's optimal amenity choice $\pi_{a,g}^*(c_{a,g}^{\pi,0})$ for every market a as a function of only the heterogeneous amenity cost shifter $c_{a,g}^{\pi,0}$. Obviously, $\partial\pi^*/\partial c_{a,g}^{\pi,0} < 0$ by the chain rule. The optimal wage is then chosen to deliver the remainder of flow utility x .⁴⁷

Since $\tilde{c}_a^\pi(0) = 0$ and $\partial\tilde{c}_a^\pi/\partial\pi(0) = 0$, a firm will always produce some amount of amenities $\pi_{a,g} > 0$. Finally, we have

$$c_{a,g}^\pi(\pi^*) = \int_0^{\pi^*} \frac{\partial\tilde{c}_a^\pi(\pi)}{\partial\pi} d\pi < \int_0^{\pi^*} 1 d\pi = \pi^*.$$

□

C.2.2 Proof of Lemma 2 (Optimal Market Selection)

Proof. Recall that composite productivity is defined as output value plus amenity value net of amenity production costs minus employer distaste, $\tilde{p}_{a,g} = pa + \pi_{a,g} - c_{a,g}^\pi(\pi_{a,g}) - z_{a,g}$. Since the vacancy cost function satisfies $c_{a,g}^v(0) = 0$ and $\partial c_{a,g}^v(0)/\partial v = 0$, a firm makes positive profits if and only if it makes positive profits per worker: $\tilde{p}_{a,g} - x > 0$. To attract workers in a market, a firm has to offer flow utility equal to or higher than the outside option of workers through a combination of wages and amenities: $x \geq \phi_{a,g}$. Therefore, a firm is profitably active in a market if and only if $\tilde{p}_{a,g} > \phi_{a,g}$. □

C.2.3 Proof of Lemma 3 (Optimal Vacancies)

Proof. We first reformulate the firm's problem. Expected profits per worker contacted by a firm is

$$\pi_{a,g}(\tilde{p}, x) = h_{a,g}(x) J_{a,g}(\tilde{p}, x),$$

where $h_{a,g}(x)$ is the acceptance probability and $J_{a,g}(\tilde{p}, x)$ is the value of employing a worker to a firm with composite productivity \tilde{p} providing flow utility x . Under the assumption that firms maximize long-run profits, the value of employing a worker is simply

$$\begin{aligned} J_{a,g}(\tilde{p}, x) &= \frac{\tilde{p} - x}{\delta_{a,g} + \lambda_{a,g}^e(1 - F_{a,g}(x)) + \lambda_{a,g}^G} \\ &= \frac{(\tilde{p} - x) / (\delta_{a,g} + \lambda_{a,g}^G)}{1 + \kappa_{a,g}^e(1 - F_{a,g}(x))}, \end{aligned}$$

⁴⁷Taking into account possible corner solutions, the optimal wage-amenity combination takes the following form:

$$\pi_{a,g}^{**}(x, c_{a,g}^{\pi,0}) = \begin{cases} x & \text{if } x < \bar{x}(c_{a,g}^{\pi,0}) \\ \pi_{a,g}^*(c_{a,g}^{\pi,0}) & \text{if } x \geq \bar{x}(c_{a,g}^{\pi,0}) \end{cases}, \quad w_{a,g}^{**}(x, c_{a,g}^{\pi,0}) = \begin{cases} 0 & \text{if } x < \bar{x}(c_{a,g}^{\pi,0}) \\ x - \pi_{a,g}^{**}(c_{a,g}^{\pi,0}, x) & \text{if } x \geq \bar{x}(c_{a,g}^{\pi,0}) \end{cases},$$

where $\bar{x}(c_{a,g}^{\pi,0})$ solves $\partial c_{a,g}^\pi(\bar{x}(c_{a,g}^{\pi,0}))/\partial\pi = 1$. Note, however, that in such corner solutions the optimal wage is $w^{**} = 0$, which is empirically not relevant. Naturally, going forward we focus on the case of an interior solution.

where $\kappa_{a,g}^e = \lambda_{a,g}^e / (\delta_{a,g} + \lambda_{a,g}^G)$. The acceptance probability for a firm offering x is

$$\begin{aligned}
h_{a,g}(x) &= \frac{u_{a,g} + s_{a,g}^e (1 - u_{a,g}) G_{a,g}(x) + s_{a,g}^G}{u_{a,g} + s_{a,g}^e (1 - u_{a,g}) + s_{a,g}^G} \\
&= \frac{\delta_{a,g} + s_{a,g}^e (\lambda_{a,g}^u + \lambda_{a,g}^G) G_{a,g}(x) + s_{a,g}^G (\delta_{a,g} + \lambda_{a,g}^u + \lambda_{a,g}^G)}{\delta_{a,g} + s_{a,g}^e (\lambda_{a,g}^u + \lambda_{a,g}^G) + s_{a,g}^G (\delta_{a,g} + \lambda_{a,g}^u + \lambda_{a,g}^G)} \\
&= \frac{1 + s_{a,g}^e \kappa_{a,g}^u G_{a,g}(x) + s_{a,g}^G (1 + \kappa_{a,g}^u)}{1 + s_{a,g}^e \kappa_{a,g}^u + s_{a,g}^G (1 + \kappa_{a,g}^u)} \\
&= \frac{1 + s_{a,g}^e \kappa_{a,g}^u \left[\frac{F_{a,g}(x)}{1 + \kappa_{a,g}^e [1 - F_{a,g}(x)]} \right] + s_{a,g}^G (1 + \kappa_{a,g}^u)}{1 + s_{a,g}^e \kappa_{a,g}^u + s_{a,g}^G (1 + \kappa_{a,g}^u)} \\
&= \frac{1 + \kappa_{a,g}^e [1 - F_{a,g}(x)] + s_{a,g}^e \kappa_{a,g}^u F_{a,g}(x) + s_{a,g}^G (1 + \kappa_{a,g}^u) [1 + \kappa_{a,g}^e [1 - F_{a,g}(x)]]}{[1 + s_{a,g}^e \kappa_{a,g}^u + s_{a,g}^G (1 + \kappa_{a,g}^u)] [1 + \kappa_{a,g}^e [1 - F_{a,g}(x)]]},
\end{aligned}$$

where $\kappa_{a,g}^u = (\lambda_{a,g}^u + \lambda_{a,g}^G) / \delta_{a,g}$ and $u_{a,g}$ is substituted with its expression in equation (8). Combining expressions, expected profits per contacted worker are

$$\begin{aligned}
\pi(\tilde{p}, x) &= h(x) J(\tilde{p}, x) \\
&= \frac{\left\{ 1 + \kappa_{a,g}^e [1 - F_{a,g}(x)] + s_{a,g}^e \kappa_{a,g}^u F_{a,g}(x) + s_{a,g}^G (1 + \kappa_{a,g}^u) [1 + \kappa_{a,g}^e [1 - F_{a,g}(x)]] \right\} (\tilde{p} - x)}{[1 + s_{a,g}^e \kappa_{a,g}^u + s_{a,g}^G (1 + \kappa_{a,g}^u)] [1 + \kappa_{a,g}^e (1 - F_{a,g}(x))]^2 (\delta_{a,g} + \lambda_{a,g}^G)}. \quad (\text{C.3})
\end{aligned}$$

Then, the firm's problem becomes

$$\max_{x,v} \left\{ \pi_{a,g}(\tilde{p}, x) v q_{a,g} - c_{a,g}^v(v) \right\},$$

where $q_{a,g}$ is defined as in equation (13). Therefore, the optimal flow-utility and vacancy policy functions satisfy

$$\begin{aligned}
x_{a,g}^*(\tilde{p}, \cdot) &= \arg \max_x \pi_{a,g}(\tilde{p}, x) \\
\frac{\partial c_{a,g}^v(v^*(\tilde{p}, \cdot))}{\partial v} &= \max_x \pi_{a,g}(\tilde{p}, x). \quad (\text{C.4})
\end{aligned}$$

Since the vacancy cost function $c^v(\cdot)$ is convex, and $\pi(\tilde{p}, x)$ in equation (C.3) is strictly increasing in \tilde{p} , then it follows from an application of the envelope theorem to equation (C.4) that $v^*(\tilde{p}, \cdot)$ is strictly increasing in \tilde{p} . Therefore, $v_{a,g}^*(\cdot)$ is strictly increasing in productivity p and strictly decreasing (constant) in z_a for women (men). Since $c_{a,g}^v(v_{a,g}) = c_{a,g}^{v,0} \times \tilde{c}^v(\pi_{a,g})$, equation (C.4) also yields that the optimal mass of vacancies is strictly decreasing in the vacancy cost shifter $c_{a,g}^{v,0}$. \square

C.2.4 Proof of Lemma 4 (Optimal Flow Utility and Wages)

Proof. We proceed in two steps.

Step 1. In the first step, we prove monotonicity of $x_{a,g}^*$ in components of $\tilde{p}_{a,g}$. Lemma 1 implies that at the optimum, amenities can be equivalently considered exogenous. Thus, we rewrite the FOCs as functions of exogenous parameters, the endogenous offer distribution, and $x_{a,g}$:

$$[\partial x_{a,g}]: \quad 1 = (\tilde{p}_{a,g} - x_{a,g}) \frac{2\lambda_{a,g}^e f_{a,g}(x_{a,g})}{\delta_{a,g} + \lambda_{a,g}^G + \lambda_{a,g}^e (1 - F_{a,g}(x_{a,g}))} \quad (\text{C.5})$$

$$[\partial v_{a,g}]: \quad c_{a,g}^{v,0} \frac{\partial \tilde{c}^v(v_{a,g})}{\partial v_{a,g}} = T_{a,g}(\tilde{p}_{a,g} - x_{a,g}) \left(\frac{1}{\delta_{a,g} + \lambda_{a,g}^G + \lambda_{a,g}^e (1 - F_{a,g}(x_{a,g}))} \right)^2, \quad (\text{C.6})$$

where $T_{a,g} = \mu_{a,g}[(u_{a,g} + s_{a,g}^G)\lambda_{a,g}^u(\delta_{a,g} + \lambda_{a,g}^G + \lambda_{a,g}^e)]/V_{a,g}$. Equation (C.5) already shows that the optimal flow utility $x_{a,g}$ is independent of the cost of posting vacancies, proving the first statement. Now consider equation (C.6); because the term on the right-hand side is always positive for $\tilde{p}_{a,g} > \phi_{a,g}$, it follows that optimal vacancies $v_{a,g}^*(\tilde{p}_{a,g}, c_{a,g}^{v,0})$ are always strictly positive.

We now show that the derivative of wages with respect to $\tilde{p}_{a,g}$ is always positive. Define $h_{a,g}(\tilde{p}_{a,g}) = F_{a,g}(x_{a,g}^*(\tilde{p}_{a,g}))$. Thus

$$h_{a,g}(\tilde{p}_{a,g}) = \frac{\int_{\tilde{p}' \geq \phi_{a,g}}^{\tilde{p}_{a,g}} \bar{v}_{a,g}^*(\tilde{p}_{a,g}) \gamma_{a,g}(\tilde{p}_{a,g})}{V_{a,g}} d\tilde{p}' \quad (\text{C.7})$$

$$h'_{a,g}(\tilde{p}_{a,g}) = f_{a,g}(x_{a,g}^*(\tilde{p}_{a,g})) x_{a,g}^{*'}(\tilde{p}_{a,g}) \quad (\text{C.8})$$

$$f_{a,g}(x_{a,g}^*(\tilde{p}_{a,g})) = h'_{a,g}(\tilde{p}_{a,g}) / x_{a,g}^{*'}(\tilde{p}_{a,g}), \quad (\text{C.9})$$

where $\bar{v}_{a,g}^*(\tilde{p}_{a,g}) = \int v_{a,g}^*(\tilde{p}_{a,g}, c') \gamma_{a,g}^c(c' | \tilde{p}_{a,g}) dc'$ is the integral of optimal vacancies conditional on $\tilde{p}_{a,g}$, $\gamma_{a,g}^c(c | \tilde{p}_{a,g})$ is the density of vacancy posting costs $c_{a,g}^{v,0}$ conditional on $\tilde{p}_{a,g}$, $\gamma_{a,g}(\tilde{p}_{a,g})$ is the marginal density of composite productivity $\tilde{p}_{a,g}$, and $\partial x_{a,g}^*(\tilde{p}_{a,g}) / \partial \tilde{p}_{a,g} = x_{a,g}^{*'}(\tilde{p}_{a,g})$ is the derivative of equilibrium flow utility with respect to $\tilde{p}_{a,g}$. Thus, we can rewrite $h'_{a,g}(\tilde{p}_{a,g}) = \frac{\bar{v}_{a,g}^*(\tilde{p}_{a,g})}{V_{a,g}} \gamma(\tilde{p}_{a,g})$ by differentiating equation (C.7) using Leibniz's integral rule.

Using these identities, we can write $f_{a,g}(x_{a,g}^*(\tilde{p}_{a,g})) = \frac{\bar{v}_{a,g}^*(\tilde{p}_{a,g})}{V_{a,g}} \gamma_{a,g}(\tilde{p}_{a,g}) \partial \tilde{p}_{a,g} / \partial x_{a,g}^*(\tilde{p}_{a,g})$. Thus, we can rewrite equation (C.5) as

$$\frac{\partial x_{a,g}^*(\tilde{p}_{a,g})}{\partial \tilde{p}_{a,g}} = (\tilde{p}_{a,g} - x_{a,g}^*) \frac{2\lambda_{a,g}^e}{\delta_{a,g} + \lambda_{a,g}^G + \lambda_{a,g}^e (1 - h_{a,g}(\tilde{p}_{a,g}))} \frac{\bar{v}_{a,g}^*(\tilde{p}_{a,g})}{V_{a,g}} \gamma_{a,g}(\tilde{p}_{a,g}). \quad (\text{C.10})$$

Because the right-hand side of this expression is always positive for $\tilde{p}_{a,g} > \phi_{a,g}$, it follows that $\partial x_{a,g}^*(\tilde{p}_{a,g}) / \partial \tilde{p}_{a,g} > 0$, thus proving that equilibrium flow utility is increasing in $\tilde{p}_{a,g}$.

Since $\tilde{p}_{a,g}$ is increasing in p and decreasing (constant) in z_a for women (men), it follows that optimal flow utility is increasing in p and decreasing (constant) in z_a for women (men).

Step 2. In the second step, we prove monotonicity of $w_{a,g}$ in components of $\tilde{p}_{a,g}$. The characterization of $w_{a,g} = x_{a,g} - \pi_{a,g}$ follows from combining Lemmas 1 and 4. \square

C.2.5 Proof of Lemma 5 (Optimal Employment)

Proof. Consider two otherwise identical employers with composite productivities $\tilde{p}_2 > \tilde{p}_1$ and optimal flow-utility and amenity choices (x_2, v_2) and (x_1, v_1) , respectively. Using the notation from

equation (15), we can write

$$\begin{aligned} & [\tilde{p}_2 - x_2] l_{a,g}(x_2, v_2) - c_{a,g}^v(v_2) > [\tilde{p}_2 - x_1] l_{a,g}(x_1, v_1) - c_{a,g}^v(v_1) \\ & > [\tilde{p}_1 - x_1] l_{a,g}(x_1, v_1) - c_{a,g}^v(v_1) > [\tilde{p}_1 - x_2] l_{a,g}(x_2, v_2) - c_{a,g}^v(v_2), \end{aligned}$$

where the first and third strict inequalities follow from uniqueness of the profit-maximizing wage choice given that firm types are distributed continuously, while the second strict inequality follows trivially. Subtracting the fourth term from the first and the third term from the second, we have $l_{a,g}(x_2, v_2) > l_{a,g}(x_1, v_1)$. This proves the comparative statics with respect to firm productivity p and the gender wedge $z_{a,g}$. The proof for the comparative statics with respect to the vacancy cost shifter $c_{a,g}^{v,0}$ is a direct consequence of the two results that the vacancy policy $v_{a,g}(\cdot)$ is strictly decreasing in $c_{a,g}^{v,0}$ (Lemma 3), while the flow-utility policy $x_{a,g}(\cdot)$ is constant in $c_{a,g}^{v,0}$ (Lemma 4). \square

C.2.6 Proof of Proposition 1 (Equilibrium Wage Equation)

Proof. We proceed in two steps. We first prove the proposition under exogenous firm-level vacancies that are constant within but may differ across genders. We then prove that under the maintained assumptions, the proposition also holds under endogenous vacancy posting.

Step 1. Suppose firms differ in their exogenous number of vacancies for each gender, $\{v_g\}_g$. Define $T_{a,g} = \mu_{a,g}[(u_{a,g} + s_{a,g}^G)\lambda_{a,g}^u(\delta_{a,g} + \lambda_{a,g}^G + \lambda_{a,g}^e)]/V_{a,g}$. First of all, Assumption 2 implies that $T_{a,g} = T_g$ for all a . Second, under exogenous vacancies, a firm's type is defined by its composite productivity $\tilde{p}_{a,g}$ and its exogenous vacancies v_g , which are constant across ability markets. As a consequence, $V_{a,g} = V_g$ in all a -markets. Using equation (14), the firm's problem can be written as

$$x_{a,g}^*(\tilde{p}_{a,g}) = \arg \max_x (\tilde{p}_{a,g} - x) \left(\frac{1}{\delta_{a,g} + \lambda_{a,g}^G + \lambda_{a,g}^e(1 - F_{a,g}(x))} \right)^2 v_g T_g.$$

Thus, given fixed vacancies, equilibrium firm profits in equation 15 can be written as

$$\Pi_{a,g}(\tilde{p}_{a,g}, v_g) = (\tilde{p}_{a,g} - x^*(\tilde{p}_{a,g})) \left(\frac{1}{\delta_{a,g} + \lambda_{a,g}^G + \lambda_{a,g}^e(1 - F_{a,g}(x^*(\tilde{p}_{a,g})))} \right)^2 v_g T_g. \quad (\text{C.11})$$

We can write the offer distribution as

$$F_{a,g}(x^*(\tilde{p}_{a,g})) = h_{a,g}(\tilde{p}_{a,g}) = \frac{1}{V_{a,g}} \int_{p' > \phi_{a,g}}^{\tilde{p}_{a,g}} \int v_g \gamma(p', v') dp' dv',$$

where $\gamma_{a,g}(p', v')$ is the joint density function of $\tilde{p}_{a,g}$ and v_g , and $h_{a,g}(\tilde{p}_{a,g})$ is the CDF of the marginal distribution of $\tilde{p}_{a,g}$, in which values are weighted by vacancies posted by firms with each particular $\tilde{p}_{a,g}$. The expression for $h_{a,g}$ is equivalent to equation (C.7) in Lemma 4, except that here we are integrating over exogenous rather than endogenous vacancies. Applying the Envelope Theorem yields

$$\frac{\partial \Pi_{a,g}(\tilde{p}_{a,g}, v_g)}{\partial \tilde{p}_{a,g}} = \left(\frac{1}{\delta_{a,g} + \lambda_{a,g}^G + \lambda_{a,g}^e(1 - h_{a,g}(\tilde{p}))} \right)^2 v_g T_{a,g} \quad (\text{C.12})$$

When $\tilde{p}_{a,g} = \phi_{a,g}$, $\Pi(\phi_{a,g}, v_g) = 0$ for all v_g , which gives us a boundary condition to solve the differential equation for profits. Rearranging (C.11) and integrating equation (C.12) yields

$$x_{a,g}(\tilde{p}_{a,g}) = \tilde{p}_{a,g} - \int_{y \geq \phi_{a,g}}^{\tilde{p}_{a,g}} \left[\frac{1 + \kappa_{a,g}^e (1 - h_{a,g}(\tilde{p}_{a,g}))}{1 + \kappa_{a,g}^e (1 - h_{a,g}(y))} \right]^2 dy, \quad (\text{C.13})$$

where $\kappa_{a,g}^e = \lambda_{a,g}^e / (\delta_{a,g} + \lambda_{a,g}^G)$. This equation parallels equation (47) in [Burdett and Mortensen \(1998\)](#), where composite productivity $\tilde{p}_{a,g}$ in the current model plays the role of job productivity differentials in their model.

Equation (C.2) and Assumption 3 imply that optimal amenities satisfy

$$\begin{aligned} c_{a,g}^{\pi,0} \times \frac{\partial \tilde{c}_a^\pi(\pi_{a,g}^*)}{\partial \pi} &= 1 \\ \frac{ac_g^{\pi,0}}{a} \frac{\partial \tilde{c}_a^\pi(\frac{\pi_{a,g}^*}{a})}{\partial \pi} &= 1, \end{aligned}$$

which obviously implies that $\pi_g^* = \pi_{a,g}^*/a$ is constant for all ability types a . Thus, amenities $\pi_{a,g}^*$ are proportional to ability. Also, the cost function implies that $c_{a,g}^\pi(\pi_{a,g}^*) = ac_g^\pi(\pi_g)$, so that also the equilibrium amenity cost is proportional to ability.

Summing up, under Assumptions 1 and 3, it follows that $\pi_{a,g} = a\pi_g$, $c_{a,g}^\pi(\pi_{a,g}) = ac_g^\pi(\pi_g)$ and $z_a = az$. Therefore, composite productivity $\tilde{p}_{a,g} = pa + \pi_{a,g} - c_{a,g}^\pi(\pi_{a,g}) - \mathbf{1}(g = F)z_a$ is proportional to a , and we can write $\tilde{p}_{a,g} = a\tilde{p}_g$, where $\tilde{p}_g = p + \pi_g - c_g^\pi(\pi_g) - \mathbf{1}(g = F)z$ is distributed according to $h_g(\tilde{p}_g)$. By definition, $h_{a,g}(\tilde{p}_{a,g}) = h_{a,g}(a\tilde{p}_g) = h_g(\tilde{p}_g)$. Owing to Assumption 2, $\kappa_{a,g} = \kappa_g$ for all a . Thus, with a change of variables and using that vacancies of each firm are constant across ability markets, we can rewrite equation (C.13) as

$$x_g(a, \tilde{p}_g) = a\tilde{p}_g - \int_{y \geq \phi_{a,g}}^{\tilde{p}_g} a \left[\frac{1 + \kappa_g^e (1 - h_g(\tilde{p}_g))}{1 + \kappa_g^e (1 - h_g(y))} \right]^2 dy. \quad (\text{C.14})$$

We still need to prove that $\phi_{a,g}$ is also proportional to a under the assumption that $b_{a,g} = ab_g$. We use a guess-and-verify approach: we guess that the case in which $\phi_{a,g}$ and equilibrium flow utility $x(\tilde{p}_g, v_g)$ are proportional to a is an equilibrium of the model and we verify it below. From equation (7) and Assumptions 1–4, we have

$$\phi_{a,g} = ab_g + (\lambda_g^u - \lambda_g^e) \int_{x' \geq \phi_{a,g}}^{\tilde{p}_g} \frac{1 - F_{a,g}(x')}{\rho + \delta_g + \lambda_g^G + \lambda_g^e (1 - F_{a,g}(x'))} dx'.$$

We proceed to show that if $\phi_{a,g} = a\phi_g$, then $x(\tilde{p}_g)$ is also proportional to a . The proof follows trivially from equation (C.14): if $\phi_{a,g} = a\phi_g$,

$$x_g(a, \tilde{p}_g) = a\tilde{p}_g - a \int_{y \geq \phi_g}^{\tilde{p}_g} \left[\frac{1 + \kappa_g^e (1 - h_g(\tilde{p}_g))}{1 + \kappa_g^e (1 - h_g(y))} \right]^2 dy.$$

Next we show that if $x(a, \tilde{p}_g)$ is proportional to a , then $\phi_{a,g}$ must be proportional to a . Consider the

bijection mapping $\tilde{p}_g(x, a) = [x^*(a, \tilde{p}_g)]^{-1}$. We can rewrite the outside option as

$$\begin{aligned}\phi_{a,g} &= ab_g + (\lambda_g^u - \lambda_g^e) \int_{x' \geq \phi_{a,g}} \frac{1 - h_{a,g}([x'(a, \tilde{p}_g)]^{-1})}{\rho + \delta_g + \lambda_g^G + \lambda_g^e(1 - h_{a,g}([x'(a, \tilde{p}_g)]^{-1}))} dx' \\ &= ab_g + a(\lambda_g^u - \lambda_g^e) \int_{x' \geq \phi_{a,g}} \frac{1 - h_g([x'(1, \tilde{p}_g)]^{-1})}{\rho + \delta_g + \lambda_g^G + \lambda_g^e(1 - h_{a,g}([x'(1, \tilde{p}_g)]^{-1}))} dx',\end{aligned}$$

which implies that the only solution to this equation satisfies $\phi_{a,g} = a\phi_g$.

Finally, recalling that $\tilde{p}_g = p + \pi_g - c_g^\pi(\pi_g) - z$ and that $w = x - \pi$, we can write monetary wages as

$$w(a, \tilde{p}_g, c_g^{\pi,0}) = a \left[\tilde{p}_g - \pi_g(c_g^{\pi,0}) - \int_{\tilde{p}' \geq \phi_g} \left[\frac{1 + \kappa_g^e(1 - h_g(\tilde{p}_g))}{1 + \kappa_g^e(1 - h_g(\tilde{p}'))} \right]^2 d\tilde{p}' \right], \quad (\text{C.15})$$

which completes the proof that the desired equilibrium wage equation holds under exogenous vacancies that are constant across ability levels.

Step 2. All that remains to be shown for the desired result to follow is that in the model with endogenous vacancy posting, we have $v_{a,g}^* = v_g^*$ for all a , so that the offer distribution $h_{a,g}$ is the same across all ability markets. Under Assumption 1, we have $c_{a,g}^{v,0} = ac_g^{v,0}$. Next, we follow a guess-and-verify approach. Suppose that $x_{a,g}^*(\tilde{p}_g)$ is proportional to ability a . Using that $F_{a,g}(x_{a,g}^*(\tilde{p}_{a,g})) = h_{a,g}(\tilde{p}_{a,g})$, we can write the first-order condition for vacancy creation in equation (C.6) as

$$\begin{aligned}ac_g^{v,0} \frac{\partial \tilde{c}_g^v(v_{a,g})}{\partial v_{a,g}} &= T_{a,g}(\tilde{p}_{a,g} - x_{a,g}^*(\tilde{p}_{a,g})) \left(\frac{1}{\delta_{a,g} + \lambda_{a,g}^G + \lambda_{a,g}^e(1 - h_{a,g}(\tilde{p}_{a,g}))} \right)^2 \\ c_g^{v,0} \frac{\partial \tilde{c}_g^v(v_{a,g})}{\partial v_{a,g}} &= T_g(\tilde{p}_g - x_g^*(\tilde{p}_g)) \left(\frac{1}{\delta_g + \lambda_g^G + \lambda_g^e(1 - h_g(\tilde{p}_g))} \right)^2,\end{aligned}$$

immediately proving that $v_{a,g} = v_g$ for all a . Equation (11) thus implies that aggregate vacancies satisfy $V_{a,g} = V_g$, which, together with Assumption 2, also implies that in equilibrium, $u_{a,g} = u_g$ and $\lambda_{a,g}^u = \lambda_g^u$. As a consequence, all terms in the wage equation (C.15) scale linearly in ability. Therefore, log wages take the form of the desired equilibrium wage equation. \square

C.3 Log-Additivity in Worker Ability

We here demonstrate log-additivity in worker ability of some key model objects—workers' value functions, workers' policy function, and firms' value function—under assumptions 1–4 of the paper.

Conjecture. We conjecture and verify that wages, amenities, and vacancies take the following form:

$$\begin{aligned}w_{g,a} &= w_g a \\ \pi_{g,a} &= \pi_g a \\ v_{g,a} &= v_g.\end{aligned}$$

Flow utility, offer distribution, and firm sizes. Defining $x_g = w_g + \pi_g$, flow utility becomes

$$x_{g,a} = x_g a.$$

The above conjecture, together with assumptions 1 and 2, implies that the offer distribution becomes

$$F_{g,a}(x_{g,a}) = F_g(x_g).$$

Given this, employment becomes

$$l_{g,a}(w_{g,a}, \pi_{g,a}, v_{g,a}) = l_g(w_g, \pi_g, v_g),$$

which allows for sorting between genders but rules out sorting within gender across ability types.

Workers' value functions. We guess and verify the following value functions:

$$\begin{aligned} S_{g,a}(xa) &= S_g(x) a \\ W_{g,a} &= W_g a \end{aligned}$$

With this guess, the value of employed workers becomes

$$\rho S_g(x) = x + \lambda_g^E \int_{x' \geq x} [S_g(x') - S_g(x)] dF_g(x') + \lambda_g^G \int_{x'} [S_g(x') - S_g(x)] dF_g(x') + \delta_g [W_g - S_g(x)]. \quad (\text{C.16})$$

Using Assumption 4, the value of unemployed workers becomes

$$\rho W_g = b_g + (\lambda_g^U + \lambda_g^G) \int_{x'} \max \{S_g(x') - W_g, 0\} dF_g(x'). \quad (\text{C.17})$$

Workers' policy function. We guess and verify the following acceptance threshold for workers' policy function:

$$\phi_{g,a} = \phi_g a.$$

With this guess, the acceptance threshold for workers' policy function becomes

$$\phi_g = b_g + (\lambda_g^U - \lambda_g^E) \int_{x' \geq \phi_g} \frac{1 - F_g(x')}{\rho + \delta_g + \lambda_g^G + \lambda_g^E [1 - F_g(x')]} dx'. \quad (\text{C.18})$$

Firms' value function. We guess and verify the following firm value function within a market:

$$\Pi_{g,a}(p, z_{g,a}, c_{g,a}^{v,0}, c_{g,a}^{\pi,0}) = \Pi_g(p, z_g, c_g^{v,0}, c_g^{\pi,0}) a.$$

Combining this guess with assumptions 3 and 4, the within-market value of a firm becomes

$$\Pi_g(p, z_g, c_g^{v,0}, c_g^{\pi,0}) = \left[p - w_g - c_g^{\pi,0} \tilde{c}(\pi_g) - z_g \right] l_g(w_g, \pi_g, v_g) - c_g^{v,0} \tilde{c}^v(v_g). \quad (\text{C.19})$$

Verifying the conjecture. That the conjecture regarding wages, amenities, and vacancies holds follows from the above results and the firm optimality conditions outlined in the paper.

Summary. We have demonstrated that under assumptions 1–4, both the equilibrium wage equation and other key model objects—workers’ value functions, workers’ policy function, and firms’ value function—are log-additive in worker ability. This implies that the labor supply curves faced by firms across worker ability markets are parallel, ability-scaled copies of one another. In other words, under the above assumptions, the economy is invariant to scale.

C.4 Alternative Modeling Assumption on Amenity Production

We here present an alternative formulation of firm’s amenity production technology. While in the baseline formulation of the model firms produce gender-specific amenity values for each worker, in the alternative formulation firms produce a vector of amenities with gender-specific utility weights for each worker. We establish conditions for *observational equivalence* and *counterfactual equivalence* between the baseline model and the alternative model. For ease of exposition, we work with the piece-rate version of our model in which wages and amenity values scale with worker ability a .

Each firm posts an amenity vector $\vec{\pi} = (\pi_1, \pi_2, \dots, \pi_N) \in \mathbb{R}^{N \times 1}$ subject to cost $c^\pi(\vec{\pi})$. Workers derive gender-specific utility from an amenity vector given by a preference vector $\vec{\beta}_g = (\beta_{g,1}, \beta_{g,2}, \dots, \beta_{g,N}) \in \mathbb{R}^{N \times 1}$. A worker of gender g at a firm with amenity vector $\vec{\pi}$ enjoys amenity utility $\vec{\pi}'\vec{\beta}_g$.

If $N = 1$, then $\vec{\pi} = \pi \in \mathbb{R}$ and $\vec{\beta}_g = \beta_g \in \mathbb{R}$. In this case, men and women agree on the ranking of employers in terms of their amenity values as long as $\text{sign}(\beta_M) = \text{sign}(\beta_F)$. Otherwise, if $\text{sign}(\beta_M) \neq \text{sign}(\beta_F)$, then men and women have opposite rankings of employers in terms of their amenity values. This formulation is too restrictive to match the data, which motivates the following two assumptions.

Assumption C.1. *Amenities are at least twofold:*

$$N \geq 2.$$

Assumption C.2. *The gender-specific preference vectors $\vec{\beta}_M^\pi$ for men and $\vec{\beta}_F^\pi$ for women are linearly independent:*

$$\nexists c \in \mathbb{R} \quad \text{s.t.} \quad \vec{\beta}_M = c\vec{\beta}_F.$$

Under these assumptions, the following result obtains, which helps us rationalize the data:

Lemma C.1 (Existence of amenity vector). *Suppose Assumptions C.1 and C.2 hold. Then for any duplet of gender-specific utilities (U_M^π, U_F^π) at a given employer, there exists an amenity vector $\vec{\pi} = (\pi_1, \pi_2, \dots, \pi_N)$ such that*

$$\begin{bmatrix} U_M^\pi \\ U_F^\pi \end{bmatrix}_{2 \times 1} = \begin{bmatrix} \vec{\beta}_M^\pi \\ \vec{\beta}_F^\pi \end{bmatrix}_{2 \times N} \vec{\pi}_{N \times 1}. \quad (\text{C.20})$$

Proof. This is simply a system of two linear equations in $N \geq 2$ unknowns. By linear independence of $\vec{\beta}_M$ and $\vec{\beta}_F$ due to Assumption C.2, the matrix that premultiplies $\vec{\pi}$ in equation (C.20) has full rank. Therefore, the system admits at least one solution. \square

If $N = 2$, then Lemma C.1 admits a unique amenity vector $\vec{\pi} = (\pi_1, \pi_2)$ that rationalizes any amenity-utility duplet (U_M^π, U_F^π) . If $N > 2$, then there exist multiple amenity vectors $\vec{\pi}$ that rationalize the same duplet of gender-specific utilities (U_M^π, U_F^π) , among which a profit-maximizing firm will pick the cost-minimizing one.

In addition, we make the following assumption as a natural extension to that in the baseline model:

Assumption C.3. Firms provide firm-wide amenities $\vec{\pi}$, and the cost of amenity provision is given by

$$c^\pi(\vec{\pi}, l_M, l_F) = \int_{j=0}^{l_M+l_F} \sum_{i=1}^N c_i^{\pi,0} \tilde{c}^\pi(\pi_i) dj,$$

where j indexes workers, i indexes amenities, $c_i^{\pi,0}$ is an amenity-specific cost shifter that differs across firms, and $\tilde{c}^\pi(\pi_i)$ is increasing convex such that $\tilde{c}^\pi(0) = 0$ and $\partial \tilde{c}^\pi / \partial \pi_i(0) = 0$.

Next, we show that there exist (unique) values of productivity p , the gender wedge z , vacancy cost functions $c_M^v(\cdot)$ and $c_F^v(\cdot)$; there also exist amenity cost shifters $c_i^{\pi,0}$ for $i \in \{1, \dots, N\}$ that rationalize a given vector of amenities $\vec{\pi}$ along with wages w_M and w_F and vacancies v_M and v_F as firms' equilibrium choices.

Lemma C.2 (Observational equivalence). *Suppose Assumptions C.1, C.2, and C.3 hold. Then, for a given firm-level amenity vector $\vec{\pi}$, there exists a firm-specific amenity cost function $c^\pi(\vec{\pi})$ such that $\vec{\pi}$ solves*

$$(\vec{\pi}, w_M, w_F, v_M, v_F) = \arg \max_{\vec{\pi}, \tilde{w}_M, \tilde{w}_F, \tilde{v}_M, \tilde{v}_F} \left\{ \sum_{g=M,F} \left[p - z_g - \tilde{w}_g - c^\pi(\vec{\pi}) \right] l_g(\vec{\pi}, \tilde{w}_g, \tilde{v}_g) - \sum_{g=M,F} c^v(\tilde{v}_g) \right\}$$

for some levels of firm productivity p , gender wedge z , and vacancy posting cost functions $c_g^v(\cdot)$.

Proof. The system of FOCs associated with firm optimality is the following:

$$[\partial \pi_i] : c_i^{\pi,0} \frac{\partial \tilde{c}^\pi(\pi_i)}{\partial \pi_i} = \frac{[p - w_M - c^\pi(\vec{\pi})] \frac{\partial l_M(\vec{\pi}, w_M, v_M)}{\partial \pi_i} + [p - w_F - c^\pi(\vec{\pi}) - z] \frac{\partial l_F(\vec{\pi}, w_F, v_F)}{\partial \pi_i}}{l_M(\vec{\pi}, w_M, v_M) + l_F(\vec{\pi}, w_F, v_F)} \quad (\text{C.21})$$

$$[\partial w_M] : 1 = [p - w_M - c^\pi(\vec{\pi})] \frac{\frac{\partial l_M(\vec{\pi}, w_M, v_M)}{\partial w_M}}{l_M(\vec{\pi}, w_M, v_M)} \quad (\text{C.22})$$

$$[\partial w_F] : 1 = [p - w_F - c^\pi(\vec{\pi}) - z] \frac{\frac{\partial l_F(\vec{\pi}, w_F, v_F)}{\partial w_F}}{l_F(\vec{\pi}, w_F, v_F)} \quad (\text{C.23})$$

$$[\partial v_M] : \frac{\partial c_M^v(v_M)}{v_M} = [p - w_M - c^\pi(\vec{\pi})] \frac{\partial l_M(\vec{\pi}, w_M, v_M)}{\partial v_M} \quad (\text{C.24})$$

$$[\partial v_F] : \frac{\partial c_F^v(v_F)}{v_F} = [p - w_F - c^\pi(\vec{\pi}) - z] \frac{\partial l_F(\vec{\pi}, w_F, v_F)}{\partial v_F}. \quad (\text{C.25})$$

Note that the only FOC containing the term $\partial \tilde{c}^\pi(\pi_i) / \partial \pi_i$ is equation (C.21). All other FOCs depend on the level of the amenity cost $c^\pi(\vec{\pi})$ but not its derivative. Hence, we can scale the amenity cost function $c^\pi(\vec{\pi})$, the productivity level for both genders, or the gender wedge for women to satisfy the wage FOCs in equations (C.22)–(C.23). Similarly, we can scale the vacancy cost $c^v(v_g)$ to satisfy the vacancy FOCs in equations (C.24)–(C.25). By Assumption C.3, there are N equations (i.e., FOCs with respect to π_i for $i = 1, 2, \dots, N$) with N free parameters (i.e., amenity cost shifters $c_i^{\pi,0}$ for $i = 1, 2, \dots, N$), so there exists N amenity cost shifters $c_i^{\pi,0}$ that satisfy firm optimality and rationalize $\vec{\pi}$. \square

Lemma C.2 establishes *observational equivalence* between our baseline model with gender-specific amenity values and the alternative formulation with an amenity vector and gender-specific utility weights. That is, any observed empirical pattern of employer ranks and pay differences by gender can be rationalized by either of the two models.

Next, we characterize optimal amenity provision. An argument analogous to that in the main paper shows that under assumptions C.1–C.3, optimal amenities satisfy

$$[\partial\pi_i] : c_i^{\pi,0} \times \frac{\partial \tilde{c}^\pi(\pi_i)}{\partial \pi_i} \sum_{g=M,F} l_g(x_g, v_g) = \beta_{M,i} l_M(x_M, v_M) + \beta_{F,i} l_F(x_F, v_F), \quad \forall i. \quad (\text{C.26})$$

Under these assumptions, optimal amenity provision depends on the gender composition of a firm's workforce, which varies with firm fundamentals and counterfactual policies. However, the model under these assumptions is at odds with the empirical observation that amenity quantities differ significantly across genders within employers, as demonstrated for a set of 12 amenities in Table E.2 of Appendix E.4. For example, in the data, women make up the vast majority of beneficiaries of parental leaves. Thus, assuming that the cost of amenities that are enjoyed by a subset of workers is paid also for all other workers seems inconsistent with the empirical evidence. Motivated by this observation, we consider the following alternative assumption in lieu of Assumption C.3.

Assumption C.4. *Firms provide individual-specific amenities $\{\bar{\pi}_j\}_j$ for each worker j and the cost of amenity provision given by*

$$c^\pi(\{\bar{\pi}_j\}_j) = \int_{j=0}^{l_M+l_F} \sum_{i=1}^N c_i^{\pi,0} \tilde{c}^\pi(\pi_{i,j}) dj,$$

where j indexes workers, i indexes amenities, $c_i^{\pi,0}$ is an amenity-specific cost shifter that differs across firms, and $\tilde{c}^\pi(\pi_{i,j})$ is increasing convex such that $\tilde{c}^\pi(0) = 0$ and $\partial \tilde{c}^\pi / \partial \pi_{i,j}(0) = 0$.

An argument analogous to that in Lemma C.2 establishes observational equivalence between our baseline model and the alternative model under Assumption C.4. Next, we characterize the dependence of a firm's optimal amenity choice on amenity cost shifters $c_i^{\pi,0}$ for $i \in \{1, \dots, N\}$ and other model parameters.

Lemma C.3 (Counterfactual equivalence). *Suppose Assumptions C.1, C.2, and C.4 hold. Then, a firm's optimal amenity policy $\pi_{i,j}^*(\cdot)$ for all (i, j) is strictly decreasing in its amenity cost shifter $c_i^{\pi,0}$, increasing in gender utility weights $\beta_{g,i}$ for $g = M, F$, and invariant to all other parameters.*

Proof. Based on the insight that workers care only about the flow utility of a job, we can rewrite the problem of a firm as one of choosing in each market a flow utility $x_g = w_g + \bar{\pi}'_g \vec{\beta}_g$ and vacancies v_g that solve the following problem:

$$\max_{\{x_g, v_g\}_{g=M,F}} \left\{ \sum_{g=M,F} [p - z_g] l_g(x_g, v_g) - c^x(x_M, x_F) - \sum_{g=M,F} c_g^v(v_g) \right\}, \quad \forall (g),$$

where $c^x(x_M, x_F)$ is the solution to the following cost-minimization subproblem in each market:

$$c^x(x_M, x_F) = \min_{w_M, w_F, \{\bar{\pi}_j\}_j} \left\{ w_M l_M(x_M, v_M) + w_F l_F(x_F, v_F) + \int_{j=0}^{l_M+l_F} \sum_{i=1}^N c_i^\pi(\pi_{i,j}) dj \right\} \quad (\text{C.27})$$

$$\begin{aligned} \text{s.t. } & w_g + \bar{\pi}'_{g(j)} \vec{\beta}_g = x_g \quad \forall j \\ & = \min_{\{\bar{\pi}_g\}_g} \left\{ (x_M - \bar{\pi}'_M \beta_M) l_M(x_M, v_M) + (x_F - \bar{\pi}'_F \beta_F) l_F(x_F, v_F) \right. \\ & \quad \left. + \sum_{i=1}^N [l_M(x_M, v_M) c^\pi(\pi_{M,i}) + l_F(x_F, v_F) c^\pi(\pi_{F,i})] \right\}, \end{aligned} \quad (\text{C.28})$$

where we move from equation (C.27) to equation (C.28) by using Assumption C.4. Note that the cost of amenity production $c^\pi(\cdot)$ and the marginal amenity utility $\vec{\beta}_g$ are identical for individual workers j of the same gender g , but different across genders. Thus, a firm's optimal amenity choice is to offer the same vector of amenities $\vec{\pi}_g$ to workers of the same gender and different amenities to workers of different genders. This model prediction is consistent with the salient empirical fact that many job amenities (e.g., parental leave benefits) are differentially accessed by men and women at the same employer. The associated optimality conditions for amenities production are the following:

$$[\partial\pi_{g,i}] : c_i^{\pi,0} \times \frac{\partial \tilde{c}^\pi(\pi_{g,i})}{\partial \pi_{g,i}} = \beta_{g,i}, \quad \forall (g,i). \quad (\text{C.29})$$

Equation (C.29) pins down a firm's optimal amenity choice $\pi_{g,i}^*(c_i^{\pi,0}, \beta_{g,i})$ as a function of the heterogeneous amenity cost shifter $c_i^{\pi,0}$ as well as the set of gender-specific amenity-utility weights $\beta_{g,i}$. Obviously, $\partial\pi_i^*/\partial c_i^{\pi,0} < 0$ and $\partial\pi_i^*/\partial\beta_{g,i} > 0$ by the chain rule. The optimal wage is then chosen to deliver the remainder of flow utility x_g to workers of gender $g = M, F$. \square

Lemma C.3 is a powerful result because it establishes *counterfactual equivalence* with respect to the equilibrium decomposition of the gender pay gap, which lies at the heart of our analysis in Section 7.1. Specifically, it follows from Lemma C.3 that the gender-specific amenity vector $\vec{\pi}_g^*$ is independent of productivity p , the gender-firm-level recruitment costs $c_g^{v,0}$, the gender wedge z , and other model parameters. Therefore, shutting down amenity cost differences across genders in counterfactual 1 of the main body of the paper has the same effects as equalizing gender preferences over the amenity vector in the alternative model.

C.4.1 Discussion of Model Choice

Both the baseline model with gender-specific amenity values and the alternative formulation with an amenity vector have attractive features. The alternative formulation seems realistic because it allows for a common set of amenities that are differentially accessible to men and women within a given employer.

A drawback of the alternative formulation is that the formulation with an amenity vector requires the strong assumption that we observe the full vector of amenities $\vec{\pi}$ or, alternatively, that the econometrician knows the gender-specific preference vectors $\vec{\beta}_g$ for $g = M, F$. In contrast, in the baseline model, we treat amenities as an unobserved gender-employer-specific characteristic that we estimate without further assumptions on the relevant set of amenities or gender-specific amenity preferences.

In the data, we find that around 50 percent of the estimated amenity values in our baseline model are accounted for by unobserved gender-firm-specific factors, which seems at odds with the assumption that we observe the full vector $\vec{\pi}$. Based on the observational equivalence (Lemma C.2) and counterfactual equivalence with respect to the equilibrium decomposition (Lemma C.3) of the two formulations under the stated assumptions, we adopt the baseline amenity-value formulation throughout the empirical analysis and for the equilibrium decomposition. When considering the equilibrium effects of policies, however, counterfactual equivalence between the two models does not generally hold. For this purpose, we proceed with our baseline model and consider robustness with regard to different parameterizations of the cost function.

C.5 Alternative Modeling Assumptions on Vacancy Posting

C.5.1 Model Alternative 1: Directed Vacancy Posting with Joint Cost Function

As a first alternative to the benchmark model, suppose that instead of the vacancy cost being separable across genders, we assume that the vacancy cost is a function of the total number of vacancies posted. This model has the strong prediction that any firm will employ either only men or only women, except in knife-edge cases.

Setup. Each firm posts a number $v_{a,M}$ of vacancies targeted at male workers and $v_{a,F}$ vacancies targeted at women. The total cost of posting vacancies $(v_{a,M}, v_{a,F})$ for men and women is given by $c_a^v(v_{a,M} + v_{a,F})$, where the function c_a^v retains the properties laid out in the main text: $c_a^v(0) = 0$, $\partial c_a^v(\cdot)/\partial v > 0$, $\partial^2 c_a^v(\cdot)/\partial v^2 > 0$.

Equilibrium characterization. To see that this setup implies gender segregation except in knife-edge cases, note that the firm's problem can now be written as

$$\max_{x_{a,M}, x_{a,F}, v_{a,M}, v_{a,F}} \left\{ \sum_{g=M,F} (\tilde{p}_{a,g} - x_{a,g}) l_{a,g}(x_{a,g}, v_{a,g}) - c_a^v(v_{a,M} + v_{a,F}) \right\}$$

The FOCs with respect to vacancy posting now read

$$[\partial v_{a,M}] : c_a^{v'}(v_{a,M} + v_{a,F}) = T_{a,M}(\tilde{p}_{a,M} - x_{a,M}) \left(\frac{1}{\delta_{a,M} + \lambda_{a,M}^G + \lambda_{a,M}^e(1 - F_{a,M}(x_{a,M}))} \right)^2, \quad (\text{C.30})$$

$$[\partial v_{a,F}] : c_a^{v'}(v_{a,M} + v_{a,F}) = T_{a,F}(\tilde{p}_{a,F} - x_{a,F}) \left(\frac{1}{\delta_{a,F} + \lambda_{a,F}^G + \lambda_{a,F}^e(1 - F_{a,F}(x_{a,F}))} \right)^2. \quad (\text{C.31})$$

Putting equations (C.30) and (C.31) into simple economic terms, the marginal cost of an additional vacancy (the left-hand side) is equated to the marginal benefit of an additional vacancy (the right-hand side). The latter consists of an increase in the employment of that worker type multiplied by the profits made per worker of that type, which is independent of the amount of vacancies posted. This is because wages are set according to other first-order conditions, which do not depend on the amount of vacancies posted by that firm.

Since the right-hand sides in equations (C.30) and (C.31) are generically not equal, except in knife-edge cases, it follows that both FOCs cannot hold. This means that the firm will be at a corner solution with regard to one of the two genders, and this must involve posting zero vacancies for that gender.

Empirical shortcomings. According to the above analysis, except for knife-edge cases, firms would hire only men or only women—whichever gives the highest marginal benefit to the firm. This model implication is empirically counterfactual since the vast majority of firms in the real world employ a mix of men and women.

C.5.2 Model Alternative 2: Undirected Vacancy Posting

As a second alternative to the benchmark model, suppose that instead of vacancies being directed separately to men and women, we assume that firms cannot discriminate between genders in their recruiting. While such a model can qualitatively account for dual-gender firms, it turns out that quantitatively, such a model clearly fails to replicate the empirical distribution of female employment shares across firms that we documented in Section 3.1.

Setup. Each firm posts a number v_a of gender-neutral vacancies for workers of each ability level at cost $c_a^v(v_a)$. In such a model, a firm's problem can be written as

$$\max_{x_{a,M}, x_{a,F}, v_a} \left\{ \sum_{g=M,F} (\tilde{p}_{a,g} - x_{a,g}) l_{a,g}(x_{a,g}, v_a) - c_a^v(v_a) \right\}.$$

Notice that we do not impose that firms hire both genders in each submarket: it is always possible for a firm to offer flow utility $x_{a,g} < \phi_{a,g}$ such that no worker of gender g will accept it. Consequently, while a total of V_a vacancies are posted in each submarket in the aggregate, only $V_{a,g} \leq V_a = \int v_a(\tilde{p}_{a,g}, c_{a,g}^{v,0}) d\Gamma_{a,g}(\tilde{p}_{a,g}, c_{a,g}^{v,0})$ vacancies are accepted in equilibrium by workers of type (a, g) . This implies that the number of matches produced in the labor market is given by

$$m_{a,g} = \chi_{a,g} [\mu_{a,g}(u_{a,g} + s_{a,g}^e(1 - u_{a,g}) + s_{a,g}^G)]^\alpha V_a^{1-\alpha} \frac{V_{a,g}}{V_a},$$

which already incorporates the probability that a worker of gender g will meet a vacancy that is associated with a wage below the reservation threshold, leading to a rejection. It is straightforward to show that this matching function exhibits all the properties of standard matching functions and that in particular, $f_{a,g}/q_{a,g} = V_a/[u_{a,g} + s_{a,g}^e(1 - u_{a,g}) + s_{a,g}^G]$, where $f_{a,g} = m_{a,g}/[u_{a,g} + s_{a,g}^e(1 - u_{a,g}) + s_{a,g}^G]$ is the job-finding rate per effective job searcher and $q_{a,g} = m_{a,g}/V_a$ is the vacancy yield rate.

Equilibrium characterization. The following equation represents the law of motion of firm sizes:

$$\begin{aligned} \dot{l}_{a,g}(x, v) = & -\delta_{a,g} l_{a,g}(x, v) - s_{a,g} \lambda_{a,g}^e (1 - F_{a,g}(x)) l_{a,g}(x, v) + \\ & v q_{a,g} \left[\frac{u_{a,g} + s_{a,g}^G}{u_{a,g} + s_{a,g}^G + (1 - u_{a,g}) s_{a,g}^e} + \frac{(1 - u_{a,g}) s_{a,g}^e}{u_{a,g} + s_{a,g}^G + (1 - u_{a,g}) s_{a,g}} G_{a,g}(x) \right]. \end{aligned}$$

Solving for the stationary solution,

$$l_{a,g}(x_{a,g}, v_a) = \left(\frac{1}{\delta_{a,g} + \lambda_{a,g}^G + \lambda_{a,g}^e (1 - F_{a,g}(x_{a,g}))} \right)^2 \frac{v_a}{V_a} \mu_{a,g} (u_{a,g} + s_{a,g}^G) \lambda_{a,g}^u (\delta_{a,g} + \lambda_{a,g}^G + \lambda_{a,g}^e). \quad (\text{C.32})$$

To find the firm's policy functions, define $T_{a,g} = \mu_{a,g} [u_{a,g} \lambda_{a,g}^u (\delta_{a,g} + s_{a,g} \lambda_{a,g}^u)] / V_a$ and composite productivity $\tilde{p}_{a,g} = ap + \pi_{a,g} - c_{a,g}^\pi (\pi_{a,g} - \mathbf{1}[g = F] z_a)$. we rewrite the firm's problem as a function of the steady state mass of employed workers as follows:

$$\begin{aligned} \max_{x_{a,M}, x_{a,F}, v_a} \left\{ T_{a,M} v_a (\tilde{p}_{a,M} - x_{a,M}) \left(\frac{1}{\delta_{a,M} + \lambda_{a,M}^G + \lambda_{a,M}^e (1 - F_{a,M}(x_{a,M}))} \right)^2 \right. \\ \left. + T_{a,F} v_a (\tilde{p}_{a,F} - x_{a,F}) \left(\frac{1}{\delta_{a,F} + \lambda_{a,F}^G + \lambda_{a,F}^e (1 - F_{a,F}(x_{a,F}))} \right)^2 - c_a(v_a) \right\}. \end{aligned}$$

The associated FOCs read

$$\begin{aligned}
c'(v_a) &= T_{a,M}(\tilde{p}_{a,M} - x_{a,M}) \left(\frac{1}{\delta_{a,M} + \lambda_{a,g}^G + \lambda_{a,M}^e(1 - F_{a,M}(x_{a,M}))} \right)^2 \\
&\quad + T_{a,F}(\tilde{p}_{a,F} - x_{a,F}) \left(\frac{1}{\delta_{a,F} + \lambda_{a,g}^G + \lambda_{a,F}^e(1 - F_{a,F}(x_{a,F}))} \right)^2 \\
1 &= (\tilde{p}_{a,M} - x_{a,M}) \frac{2\lambda_{a,M}^e f_{a,M}(x_{a,M})}{\delta_{a,M} + \lambda_{a,M}^G + \lambda_{a,M}^e(1 - F_{a,M}(x_{a,M}))} \\
1 &= (\tilde{p}_{a,F} - x_{a,F}) \frac{2\lambda_{a,F}^e f_{a,F}(x_{a,F})}{\delta_{a,F} + \lambda_{a,F}^G + \lambda_{a,F}^e(1 - F_{a,F}(x_{a,F}))}.
\end{aligned}$$

Empirical shortcomings. Recall from Section 3.1 that firm-level female employment shares are dispersed, ranging from almost 0 to almost 1 in the data. It is this salient feature of the data that the undirected vacancy-posting model fails to replicate. To demonstrate this, we show that analytically derived expressions for the lowest and highest female employment shares are inconsistent with the data for realistic calibrations of the labor market parameters guiding worker flows.

Using equation (C.32), we can write the female share of a firm as

$$\begin{aligned}
s_f &= \frac{l_{a,F}(x_{a,F}, v_a)}{l_{a,F}(x_{a,F}, v_a) + l_{a,M}(x_{a,M}, v_a)} \\
&= \frac{1}{1 + \frac{\left(\frac{1}{\delta_{a,M} + \lambda_{a,M}^G + \lambda_{a,M}^e(1 - F_{a,M}(x_{a,M}))} \right)^2 (u_{a,M} + s_{a,M}^G) \lambda_{a,M}^u (\delta_{a,M} + \lambda_{a,M}^G + \lambda_{a,M}^e)}{\left(\frac{1}{\delta_{a,F} + \lambda_{a,F}^G + \lambda_{a,F}^e(1 - F_{a,F}(x_{a,F}))} \right)^2 (u_{a,F} + s_{a,F}^G) \lambda_{a,F}^u (\delta_{a,F} + \lambda_{a,F}^G + \lambda_{a,F}^e)}}.
\end{aligned}$$

On the right-hand side, we can substitute our empirical estimates of the U-E transition rates $\lambda_{a,g}^u$, the E-U transition rates $\delta_{a,g}$, the compulsory offer arrival rate $\lambda_{a,g}^G$, the voluntary offer arrival rate $\lambda_{a,g}^e$, and the compulsory offer search units $s_{a,g}^G$ from Table 5. Thus, we obtain:

$$\begin{aligned}
(u_{a,M} + s_{a,M}^G) \lambda_{a,M}^u (\delta_{a,M} + \lambda_{a,M}^G + \lambda_{a,M}^e) &= (0.243 + 0.119) \times 0.100 \times (0.036 + 0.012 + 0.006) \approx 0.0020 \\
(u_{a,F} + s_{a,F}^G) \lambda_{a,F}^u (\delta_{a,F} + \lambda_{a,F}^G + \lambda_{a,F}^e) &= (0.240 + 0.107) \times 0.087 \times (0.031 + 0.009 + 0.005) \approx 0.0014.
\end{aligned}$$

Thus, this ratio is approximately equal to $0.0020/0.0014 = 1.429$, and the expression simplifies to

$$s_f = \frac{1}{1 + 1.429 \times \left(\frac{\delta_{a,F} + \lambda_{a,F}^G + \lambda_{a,F}^e(1 - F_{a,F}(x_{a,F}))}{\delta_{a,M} + \lambda_{a,M}^G + \lambda_{a,M}^e(1 - F_{a,M}(x_{a,M}))} \right)^2} = \frac{1}{1 + 1.429 \times \left(\frac{0.040 + 0.005 \times (1 - F_{a,F}(x_{a,F}))}{0.048 + 0.006 \times (1 - F_{a,M}(x_{a,M}))} \right)^2}.$$

Since firm sizes are monotonically increasing in flow utility x offered by the firm, we can obtain expressions for the minimum female employment share \underline{s}_f and the maximum female employment share \bar{s}_f by focusing on employers that are at the very top of the job ladder for one gender and simultaneously at the very bottom of the job ladder for the other gender. Specifically, among all dual-gender firms, the firm with the highest female employment share has $F_{a,F} = 1$ and $F_{a,M} = 0$. Conversely, the firm with the lowest female employment share has $F_{a,F} = 0$ and $F_{a,M} = 1$.

Thus, we find that the minimum (maximum) female employment share in the model is ≈ 0.444 (0.561), which is inconsistent with the minimum (maximum) female employment share's being close

to 0 (1) in the data.

Using the numerical solution algorithm proposed in Appendix C.5.3, we further scrutinize the properties of the undirected vacancy-posting model.

C.5.3 Solution Algorithm for Undirected-Vacancy-Posting Model

The undirected-vacancy-posting model is challenging because it does not allow us to solve men and women as two separate differential equations. Thus, we rely on a different algorithm to solve the undirected vacancy-posting model.

Assume $c(v_a) = c \frac{v_a^2}{2}$. Then,

$$v_a = \frac{T_{a,M}}{c} (\tilde{p}_{a,M} - x_{a,M}) \left(\frac{1}{\delta_{a,M} + \lambda_{a,M}^G + \lambda_{a,M}^e (1 - F_{a,M}(x_{a,M}))} \right)^2 + \frac{T_{a,F}}{c} (\tilde{p}_{a,F} - x_{a,F}) \left(\frac{1}{\delta_{a,F} + \lambda_{a,M}^G + \lambda_{a,F}^e (1 - F_{a,F}(x_{a,F}))} \right)^2.$$

By definition:

$$\begin{aligned} V_{a,M} &= \int \int v_a \mathbb{I}[\tilde{p}_{a,M} > \phi_{a,M}] \gamma(\tilde{p}_{a,M}, \tilde{p}_{a,F}) d\tilde{p}_{a,M} d\tilde{p}_{a,F} \\ V_{a,F} &= \int \int v_a \mathbb{I}[\tilde{p}_{a,F} > \phi_{a,F}] \gamma(\tilde{p}_{a,M}, \tilde{p}_{a,F}) d\tilde{p}_{a,M} d\tilde{p}_{a,F} \\ V_a &= \int \int v_a \mathbb{I}[\tilde{p}_{a,M} > \phi_{a,M} \vee \tilde{p}_{a,F} > \phi_{a,F}] \gamma(\tilde{p}_{a,M}, \tilde{p}_{a,F}) d\tilde{p}_{a,M} d\tilde{p}_{a,F}. \end{aligned}$$

Start by defining $\gamma(\tilde{p}_{a,M})$ as the marginal distribution of $\tilde{p}_{a,M}$. Also, define

$$\bar{v}_a(\tilde{p}_{a,M}) = \int \frac{v(\tilde{p}_{a,M}, \tilde{p}_{a,F}) \gamma(\tilde{p}_{a,M}, \tilde{p}_{a,F})}{\gamma(\tilde{p}_{a,M})} d\tilde{p}_{a,F}$$

as average vacancies posted by firms with composite productivity for men $\tilde{p}_{a,M}$. Then we can write

$$\begin{aligned} h(\tilde{p}) &= F(x(\tilde{p})) \\ \implies h'(\tilde{p}) &= f(x(\tilde{p})) x'(\tilde{p}) \\ \implies f(x(\tilde{p})) &= h'(\tilde{p}) / x'(\tilde{p}) \\ \bar{v}_a(\tilde{p}_{a,M}) &= \frac{V_{a,M} h'(\tilde{p}_{a,M})}{\gamma(\tilde{p}_{a,M})} \\ \implies h'(\tilde{p}_{a,M}) &= \frac{\bar{v}_a(\tilde{p}_{a,M})}{V_{a,M}} \gamma(\tilde{p}_{a,M}). \end{aligned} \tag{C.33}$$

Thus, the flow utility FOC can be rewritten as follows:

$$\begin{aligned} 1 &= (\tilde{p}_{a,g} - x_{a,M}) \frac{2\lambda_{a,g}^e f_{a,g}(x_{a,g})}{\delta_{a,g} + \lambda_{a,g}^G + \lambda_{a,g}^e (1 - F_{a,g}(x_{a,g}))} \\ x'(\tilde{p}_{a,g}) &= (\tilde{p}_{a,g} - x_{a,g}) \frac{2s_{a,g} \lambda_{a,g}^u h'(p_{a,g})}{\delta_{a,g} + \lambda_{a,g}^G + \lambda_{a,g}^e (1 - h(\tilde{p}_{a,g}))} \end{aligned} \tag{C.34}$$

We start from the case in which $\gamma(\tilde{p}_{a,M}, \tilde{p}_{a,F})$ is an analytical function and the associated marginal and conditional distributions are easy to compute. In this case, the marginals $\gamma(\tilde{p}_{a,g})$ are also known. Using the previous intuitions, the algorithm works as follows:

1. Start with a guess for $x(\tilde{p}_{a,g})$, $F(x(\tilde{p}_{a,g}))$ and $V_{a,g}$ for each gender.
2. Using the guess for $V_{a,g}$, calculate transition rates $\lambda_{a,g}^u$, $\lambda_{a,g}^e$ and $\lambda_{a,g}^G$.
3. Using the guess for the flow utility function, compute $v_a(\tilde{p}_{a,M}, \tilde{p}_{a,F})$ on a large grid over values of $\tilde{p}_{a,M}$ and $\tilde{p}_{a,F}$.
4. Normalize v_a to get $V_{a,g}$. Use v_a to compute $h'(\tilde{p}_{a,g})$ for both genders using equation (C.33).
5. Use $h'(\tilde{p}_{a,g})$ to compute the new function $h(\tilde{p}_{a,g}) = F(x(\tilde{p}_{a,g}))$.
6. Use h' and h to solve the ordinary differential equation (C.34).
7. Update the wage function $x(\tilde{p}_{a,g})$ and the CDF $F(x(\tilde{p}_{a,g}))$.
8. Go back to step 2 and repeat until convergence.


D Identification Appendix

D.1 Illustrative Identification Example

To illustrate how we estimate employer ranks, productivity, amenities, and gender wedges from data on worker flows and pay across establishments, we use a simple example. For the purpose of this simple example, we abstract from endogenous vacancy and amenity creation and heterogeneity in the offer densities and labor market parameters, all of which will be present in the general estimation routine.

Consider three employers A , B , and C and a pool of nonemployed workers N . Because the PageRank does not depend on employer size, we can think of all three employers as having a large number of male and female workers. To simplify the example, we assume that each employer hires a fixed number of male and female workers but from different sources. Employer A hires men in equal proportions by poaching from B , C , and nonemployment; employer B hires from C and nonemployment most of the time but rarely from A ; while employer C hires from nonemployment most of the time but rarely from A and B . Female worker flows are identical except that employer B hires women from nonemployment most of the time but rarely from A and C ; while employer C hires from B and nonemployment most of the time but rarely from A . Figure D.1 graphically summarizes the labor markets for men and women.

Figure D.1. Example worker flows between nonemployment and employers, by gender



Note: Figure shows example of worker flows between employers. Nodes A , B , and C represent different employers. Node N represents nonemployment. Arrows represent worker flows. Numbers above arrows represent share of all worker flows from a given node.

Estimating PageRanks based on equation (4) of Section 3.2 for this labor market yields separate employer rankings by gender. Intuitively, employers that poach a lot of workers of a given gender from other high-ranked employers are themselves highly ranked according to the PageRank. For men, employer A is ranked highest (PageRank index 0.423), employer B is middle-ranked (PageRank index 0.326), and employer C is ranked lowest (PageRank index 0.251). For women, employer A is also ranked highest (PageRank index 0.423), employer B is ranked in lowest (PageRank index 0.251), and employer C is middle-ranked (PageRank index 0.326). Men and women agree on employer A being ranked highest but disagree on the ranking of the remaining two employers B and C . The resulting PageRanks for men and women are in columns (1)–(2) of Table D.1.

Suppose that pay at employers (A, B, C) is $(8.0, 7.0, 4.0)$ for men and $(8.0, 6.0, 3.8)$ for women, as in columns (3)–(4) of Table D.1. Suppose also that the underlying amenity values at those employers are $(1.0, 0.0, 0.0)$ for men and $(0.0, 0.0, 2.3)$ for women, as in columns (5)–(6) of Table D.1. Finally, suppose that the underlying employer productivities are $(15.9, 12.3, 4.0)$, as in column (13), and gender wedges

are (5.3, 6.3, 0.0), as in column (15) of Table D.1.

Table D.1. Example pay, PageRanks, amenities, utilities, productivities, gender wedges

	PageRanks		Pay		Amenities				Utilities				Prod.		Wedge	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
Emp.	r_M	r_F	w_M	w_F	π_M	π_F	$\hat{\pi}_M$	$\hat{\pi}_F$	x_M	x_F	\hat{x}_M	\hat{x}_F	p	\hat{p}	z	\hat{z}
N	-	-	-	-	-	-	-	-	4.0	6.0	4.0	6.0	-	-	-	-
A	3	3	8.0	8.0	1.0	0.0	0.0	0.0	9.0	8.0	8.0	8.0	15.9	14.6	5.3	4.0
B	2	1	7.0	6.0	0.0	0.0	0.0	0.0	7.0	6.0	7.0	6.0	12.3	12.3	6.3	6.3
C	1	2	4.0	3.8	0.0	2.3	0.0	2.3	4.0	6.1	4.0	6.1	4.0	4.0	0.0	0.0

Note: Table shows example estimates of employer ranks using PageRank, pay, amenities, utilities, productivity, and gender wedges. "Emp." stands for employer A, B, C, or nonemployment status N. For each gender g , w_g is wage, r_g is Pagerank, π_g is amenity value, p is employer productivity ("prod."), and z is gender wedges ("wedge"). Hats denote estimates.

How can the information on (estimated) employer ranks and pay be used to infer (unobserved) employer productivity, amenities, and gender wedges? We proceed in five steps.

First, we pick gender-specific employer amenities to make employer ranks consistent with pay by gender. Without loss of generality, assume amenity values are weakly positive. For men, it must be that $\pi_M^A > -1.0 + \pi_M^B$ and $\pi_M^A > -4.0 + \pi_M^C$ for A to be highest-ranked, and $\pi_M^B > -3.0 + \pi_M^C$ for B to be middle-ranked. The amenities-minimizing estimate that satisfies these inequalities is $(\hat{\pi}_M^A, \hat{\pi}_M^B, \hat{\pi}_M^C) = (0.0, 0.0, 0.0)$. Similarly, for women, it must be that $\pi_F^A > -2.0 + \pi_F^B$ and $\pi_F^A > -4.2 + \pi_F^C$ for A to be highest-ranked and $\pi_F^B < -2.2 + \pi_F^C$ for B to be lowest-ranked. The amenities-minimizing estimate that satisfies these inequalities is $(\hat{\pi}_F^A, \hat{\pi}_F^B, \hat{\pi}_F^C) = (0.0, 0.0, 2.3)$. These estimates are summarized in columns (7)–(8) of Table D.1.

Second, we derive an estimate of the outside option value by gender. To this end, we first compute utility $\hat{x}_g^r = w_g^r + \hat{\pi}_g^r$ as the sum of pay and estimated amenity values for each employer. Based on the above estimates, we have $(\hat{x}_M^A, \hat{x}_M^B, \hat{x}_M^C) = (8.0, 7.0, 4.0)$ and $(\hat{x}_F^A, \hat{x}_F^B, \hat{x}_F^C) = (8.0, 6.0, 6.1)$, shown in columns (11)–(12) of Table D.1. Defining the outside option value as the lowest utility among employed workers for each gender, $\phi_g = \min_r x_g^r$, we have $\hat{\phi}_M = \hat{\pi}_M^N = 4.0$ and $\hat{\phi}_F = \hat{\pi}_F^N = 6.0$, as shown in the top row of columns (9)–(10) of Table D.1.

Third, we deduce composite productivities \tilde{p}_g^r for $g = M, F$ based on the equilibrium wage equation (17), which relates w_g^r to \tilde{p}_g^r , π_g^r , κ_g^e , and $F_g(x_g^r)$. Since we have already estimated $\kappa_g^e = \lambda_g^e / (\delta_g + \lambda_g^C)$ in Step 2 above, we assume $\kappa_M^e = \kappa_F^e = 1$ for this example. Approximating the integral in equation (17) by use of the lower Riemann (Darboux) sum, we have

$$\tilde{p}_g^r \approx x_g^r + \sum_{r' \leq r} \left(\frac{1 + \kappa_g^e [1 - F_g^r]}{1 + \kappa_g^e [1 - F_g^{r'-1}]} \right)^2 (\tilde{p}_g^r - \tilde{p}_g^{r-1}). \quad (\text{D.1})$$

We use equation (D.1) recursively with $(F_M^A, F_M^B, F_M^C) = (1, 2/3, 1/3)$, $(F_F^A, F_F^B, F_F^C) = (1, 1/3, 2/3)$, and $\tilde{p}_g^1 = \phi_g$ to estimate $(\hat{\tilde{p}}_M^A, \hat{\tilde{p}}_M^B, \hat{\tilde{p}}_M^C) = (14.6, 12.3, 4.0)$ and $(\hat{\tilde{p}}_F^A, \hat{\tilde{p}}_F^B, \hat{\tilde{p}}_F^C) = (10.6, 6.0, 6.3)$. Note that the estimated composite productivities satisfy monotonicity with respect to estimated utilities \hat{x}_g^r and PageRanks.

Fourth, we turn to men in order only to derive an estimate of employer productivity from the estimated composite productivity and amenity values. Since $z_M = 0$ by normalization, the definition of composite productivity for men yields $p = \tilde{p}_M^r - \pi_M^r$. The resulting productivity estimates are

$(\hat{p}^A, \hat{p}^B, \hat{p}^C) = (14.6, 12.3, 4.0)$, shown in column (14) of Table D.1.

Finally, we turn to women only in order to estimate gender wedges from the estimated composite productivity, amenity values, and productivity. The definition of composite productivity for women yields $z^r = p + \pi_F^r - \tilde{p}_F^r$. The resulting estimates of gender wedges are $(\hat{z}^A, \hat{z}^B, \hat{z}^C) = (4.0, 6.3, 0.0)$, shown in column (16) of Table D.1.

How do we interpret these results? Our estimates confirm that pay gaps are not utility gaps and that higher utilities are associated with higher composite productivity (Lemma 4). Focusing on employer A , we learn that equal pay (or, hypothetically, equal utility) across genders within an employer does not imply a zero gender wedge. This is because the gender wedge captures the degree to which an employer underpays or overpays relative to the competitive benchmark described by the equilibrium wage equation. In this case, women at employer A are paid lower relative to the value of their outside option compared with men. Focusing on employer B , we see that the gender wedge may be nonmonotonic across ranks based on revealed preference, pay, or productivity. This is because differently-ranked employers may either underpay or overpay relative to the competitive benchmark. Focusing on employer C , we note that even employers with a zero gender wedge may deliver different pay and utility to men compared with women. This is because differences in the outside option value, owing to either gender differences in the flow values of nonemployment or the presence of other employers with nonzero gender wedges, are priced into wage and utility offers in equilibrium.

It is worth noting that the parameter estimates in columns (7)–(8), (11)–(12), (14) and (16) of Table D.1 imperfectly approximate the underlying parameter values in columns (5)–(6), (9)–(10), (13) and (15). Naturally, the approximation is more precise in the middle of the employer rank distribution and becomes more accurate as we increase the number of employers in the data.⁴⁸

D.2 Further Details on Identification

A challenge in estimating productivity is that $f_g(x_g^r)$ is unknown, because it is the density function in the space of flow utilities x , rather than the change in the offer distribution f_g^r across ranks that we estimate. We begin by substituting f_g^r with the kernel density estimate \hat{f}_g^r , for computational stability. In other words, we need to transform the density through a change of variables: $\hat{f}_g(x_g^r) = \hat{f}_g^r \partial x / \partial r$. To perform the change of variables, we approximate the derivative by implicitly inserting the constraints in our algorithm. By definition, from one rank to the next, $\partial r = 1$. Then, we approximate $\partial x = x_g^{r+1} - x_g^r$ and rewrite the problem as follows:

$$\begin{aligned} \min_{\{\pi_g^1, \dots, \pi_g^{R^g}\}} \sum_r & \left[(w_g^{r+1} + \pi_g^{r+1}) - (w_g^r + \pi_g^r) \right]^2 \\ \text{s.t. } & w_g^r + \pi_g^r \leq w_g^{r+1} + \pi_g^{r+1}, \quad \forall r < R^g \\ & w_g^r + \pi_g^r + \frac{1 + \kappa_g^e (1 - F_g(x_g^r))}{2\kappa_g^e \hat{f}_g^r} \left[(w_g^{r+1} + \pi_g^{r+1}) - (w_g^r + \pi_g^r) \right] \\ & \leq w_g^{r+1} + \pi_g^{r+1} + \frac{1 + \kappa_g^e (1 - F_g(x_g^{r+1}))}{2\kappa_g^e \hat{f}_g^{r+1}} \left[(w_g^{r+2} + \pi_g^{r+2}) - (w_g^{r+1} + \pi_g^{r+1}) \right], \quad \forall r < R^g. \end{aligned}$$

After this substitution, the problem is written only as a function of known data inputs $w_g^r, F_g^r, f_g^r, \lambda_g^e, \lambda_g^G, \delta_g$ and of the unknowns of the problem $\{\pi_g^1, \dots, \pi_g^{R^g}\}$.

⁴⁸In the real data, we work with hundreds of thousands of employers for each gender.

D.3 Identifying Productivity and Amenities in Monte Carlo Simulations

In this subsection, we perform Monte Carlo simulations of our model and use our estimation algorithm to recover the underlying distribution of firm-level parameters, based only on the same information we observe in the data, as detailed in Section 5.

For the purpose of this exercise, we only need to focus on one gender, with the understanding that the simulations recover p if the algorithm is run on men’s data and $p - z$ if the algorithm is run on women’s data. We start by drawing 300,000 firms, characterized by productivity p and amenities π , jointly normally distributed with correlation ρ . We then transform productivity to have Pareto marginal distribution. Finally, we draw firm-specific vacancy posting cost shifters $c^{v,0}$ that are increasing in productivity and have a uniform random component. We do so because in the data, we find firm sizes to be less increasing in productivity than our model would otherwise predict. Thus, we choose a specification that replicates this moment of the data by setting firm-specific $c^{v,0}$ such that

$$c^{v,0} = 10^5 \times p \times (0.5 + U),$$

where $U \sim [0, 1]$ is a uniform random variable.

We then feed this nonparametric joint distribution of $\{p, \pi, c^{v,0}\}$ to our discrete solution algorithm and solve our model. The output of the simulation are firm-level wages, amenities, ranks and hiring intensities. We use this data to construct our estimates of rank r , density f_r , and CDF F_r as explained in Subsection D.2. Finally, in order to test whether our algorithm is successful at uncovering the true firm-specific parameters, we use only data on firm-level ranks r , wages w_r , densities f_r , and CDF F_r to estimate amenities and productivity at the firm level.

Our results are summarized in Table D.2, which shows moments of the distribution of recovered estimates under different parametrizations of the underlying amenities distribution. Under different parameterizations of the data-generating process, shown in columns (1)–(5) of the table, our algorithm recovers estimates of the amenity values and productivities that are statistically close to the true values. To evaluate the goodness of fit of our procedure, Panel E shows the correlations (ρ) and mean squared error (MSE) between our amenity estimates and true amenity values ($\rho(\hat{\pi}, \pi)$) and those between our productivity estimates and true productivity values ($\rho(\hat{p}, p)$). The correlation for amenities ranges from 0.920 (column [4]) to 0.997 (column [3]), while that for productivity ranges from 0.928 (column [3]) to 0.976 (column [5]) across simulations. The MSE for amenities ranges from 0.001 (column [3]) to 0.028 (column [4]), while that for productivity ranges from 0.236 (column [3]) to 1.956 (column [4]). We conclude that our estimation routine provides a good statistical fit in Monte Carlo simulations.

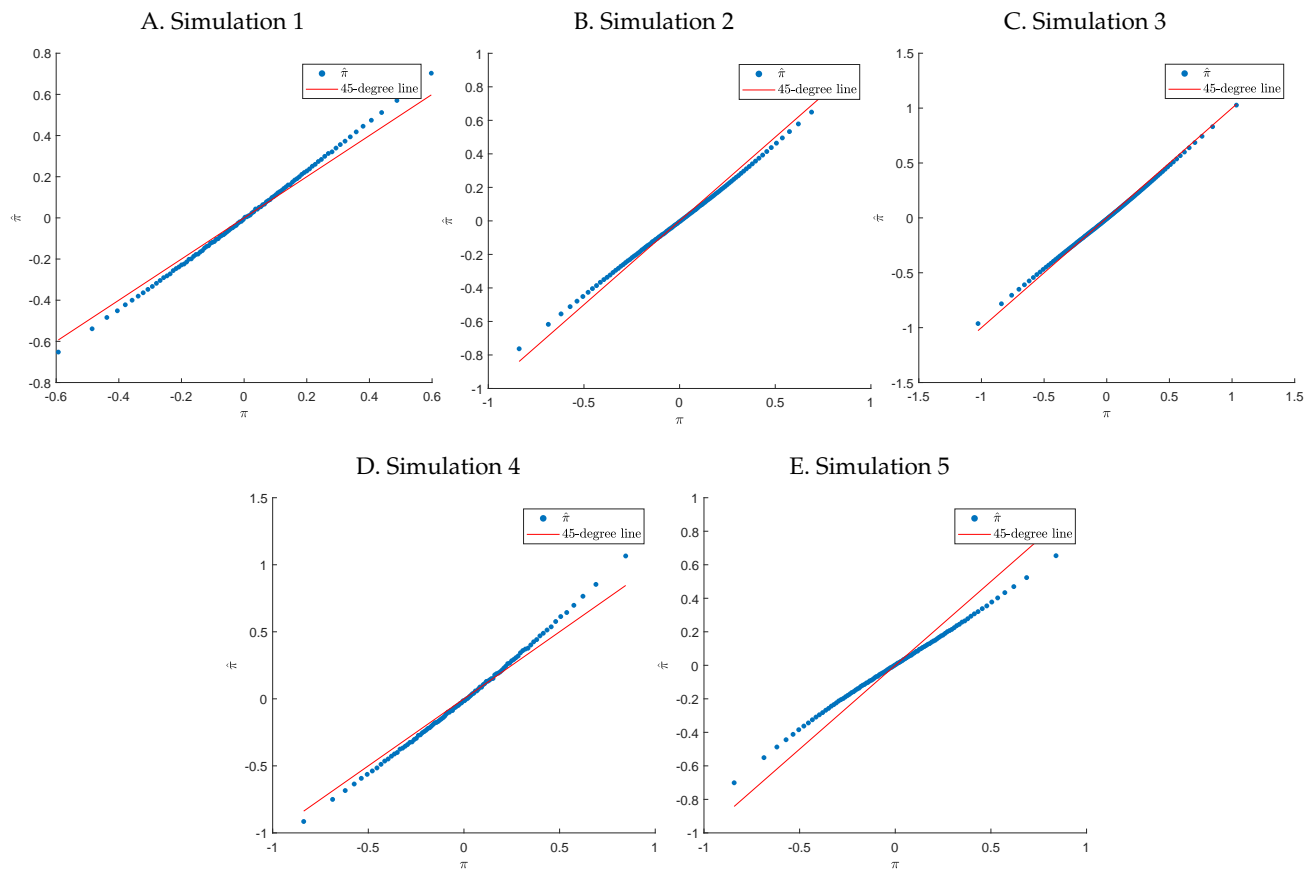
Figure D.2 visualizes the fit of our estimation routine with respect to the main object of interest—namely, the amenity values. Across all five simulations, amenity estimates are close to the 45-degree line, suggesting that our estimation routine provides a very good—albeit not perfect—fit to the simulated data.

Table D.2. Monte Carlo simulations

	(1)	(2)	(3)	(4)	(5)
<i>Panel A. Properties of true wages w</i>					
$\rho(w,r)$	0.348	0.156	0.058	0.245	0.057
<i>Panel B. Properties of true amenity values π</i>					
$\sigma(\pi)$	0.224	0.316	0.387	0.316	0.306
$\rho(\pi,w)$	-0.636	-0.698	-0.715	-0.674	-0.549
$\rho(\pi,r)$	0.479	0.583	0.644	0.538	0.783
<i>Panel C. Properties of true productivity p</i>					
$\sigma(p)$	0.511	0.514	0.514	0.510	0.511
$\rho(p,w)$	0.756	0.696	0.675	0.768	0.511
$\rho(p,r)$	0.803	0.748	0.707	0.735	0.819
$\rho(p,\pi)$	-0.003	-0.001	-0.001	-0.082	0.410
<i>Panel D. Properties of estimated amenities $\hat{\pi}$, prod. \hat{p}</i>					
$\sigma(\hat{\pi})$	0.272	0.288	0.367	0.404	0.243
$\sigma(\hat{p})$	0.844	0.874	0.848	1.507	0.471
$\rho(\hat{\pi},w)$	-0.351	-0.752	-0.729	-0.350	-0.740
$\rho(\hat{\pi},r)$	0.717	0.486	0.612	0.771	0.607
$\rho(\hat{p},w)$	0.593	0.534	0.432	0.519	0.467
$\rho(\hat{p},r)$	0.924	0.895	0.911	0.896	0.883
$\rho(\hat{p},\hat{\pi})$	0.517	0.145	0.276	0.614	0.222
<i>Panel E. Goodness of fit</i>					
$\rho(\hat{\pi},\pi)$	0.945	0.990	0.997	0.920	0.968
MSE	0.009	0.003	0.001	0.028	0.009
$\rho(\hat{p},p)$	0.962	0.950	0.928	0.931	0.976
MSE	0.411	0.285	0.236	1.956	0.441

Note: Table reports estimation results using simulated data from 300,000 firms under different parametrizations of the underlying distribution of firm heterogeneity, including amenity values π , productivity p , wages w , and employer ranks r . σ denotes standard deviation, ρ denotes correlation, and MSE denotes the mean squared error. Columns (1)–(5) show separate simulations under different true parameterizations of the data-generating process described in the text of Appendix D.3.

Figure D.2. Amenity estimates against true amenities in Monte Carlo simulations



Note: Figure plots average amenity estimates against percentile bins of true amenity values. Subfigures A–E correspond to columns (1)–(5) in Table D.2.

E Estimation Appendix

E.1 Sensitivity Analysis across Different Employer Rank Measures

To check how sensitive our results are to the choice of ranking, we re-estimate employer-specific parameters—productivity p , amenity values π_M and π_F , and gender wedges z —by applying our estimation routine to three alternative ranking measures: the Pagerank (as in the main text), the poaching index (Moscarini and Postel-Vinay, 2008; Bagger and Lentz, 2019), and the net poaching index (Haltiwanger et al., 2018; Moscarini and Postel-Vinay, 2018). Table E.1 presents employment-weighted correlations between various estimation objects across different employer rank measures. Estimates across rank measures are significantly positively, albeit not perfectly, correlated. For our baseline analysis in the main text, we use the PageRank index because it utilizes the most information per observed worker transition.

Table E.1. Counterfactual simulations, shutting down differences across gender

<i>Panel A. Productivity estimates</i>			
	Pagerank	Poaching index	Net poaching index
PageRank	1.000		
Poaching index	0.338	1.000	
Net poaching index	0.457	0.379	1.000
<i>Panel B. Male amenity estimates</i>			
	PageRank	Poaching index	Net poaching index
PageRank	1.000		
Poaching index	0.442	1.000	
Net poaching index	0.437	0.648	1.000
<i>Panel C. Female amenity estimates</i>			
	PageRank	Poaching index	Net poaching index
PageRank	1.0000		
Poaching index	0.420	1.000	
Net poaching index	0.446	0.864	1.000

Note: Table reports pairwise correlations between estimates of productivity (Panel A), male amenities (Panel B), and female amenities (Panel C) based on different employer rank measures.

E.2 Details on Covariates for Analysis of Amenity Estimates

We include as covariates in $E_{g,j}$ in equation (25) the following fourteen variables that we construct using the RAIS data: an indicator for whether the employer provides in-kind remuneration in the form of food stamps; the shares of workers with part-time contracts, with hours changes since the previous year, with paid sick leave, with parental leave, with unpaid leave, with earnings cuts since the previous year, with noncontractual earnings fluctuations, with work-related accidents, with commute-related accidents; the shares of worker separations due to firing for unjust reasons and due to worker death; five-digit industry dummies; and municipality dummies.

E.3 Details on Covariates for Analysis of Recruiting Cost Estimates

We include as covariates in E_j in equation (26) the following seven variables that we construct using the RAIS data: employer size measured as the log number of employees, an indicator for whether the highest-paid worker is a woman, the share of workers with a college degree, the share of workers with a high-school degree (the omitted category being less than a high-school degree), mean log years of worker age, mean log months of worker tenure, the share of workers with commute-related accidents, five-digit industry dummies, and municipality dummies.

E.4 Within-Employer Gender Differences in Amenity Utilization

Certain employer amenities (e.g., location, building facilities) might be equally enjoyed across genders, while others might be differentially offered to or accessed by women compared with men (e.g., parental leaves (Kleven et al., 2019), sexual harassment (Folke and Rickne, 2020)). To what extent are amenities enjoyed equally across genders within the same employer? To answer this question, we focus here on a set of 12 observable amenity proxies.⁴⁹ We project gender-employer means of these amenity proxies on a set of gender dummies while controlling for a full set of employer dummies. Table E.2 shows the results of this exercise for all dual-gender employers in our sample. We find that within the same employer, all 12 amenity proxies are accessed differentially by men and women. Women are significantly more likely to receive employer-provided food stamps; work part time; change hours; and go on sick leave, parental leave, and unpaid leave. Conversely, men are more likely to experience earnings cuts, noncontractual-earnings fluctuations, work-related and commute-related accidents, and separations due to firing for unjust reasons or worker death.

These differences are often economically meaningful relative to the baseline mean amenity proxies for men compared with women. For example, women are 3.1 percentage points, or by around two orders of magnitude, more likely to go on parental leave, and almost all of this difference is between genders within employers, with an estimated coefficient of 3.097 percent.

In summary, we find that although amenities vary significantly across employers, there are statistically and economically large differences in amenity utilization across genders within employers.

E.5 Details on Covariates of Wedge Estimates

We include as covariates in the vector E_j in equation (27) the following 12 variables, which we construct using the RAIS data: the mean intensity for routine-manual tasks, nonroutine-manual tasks, routine-cognitive tasks, nonroutine cognitive tasks involving interpersonal skills, nonroutine cognitive tasks involving analytical skills,⁵⁰ the share of worker separations due to worker death, the share of workers with work-related accidents, the female employment share, an indicator for whether the highest-paid worker is a woman, an indicator for whether an employer has no major financial stakeholder (as proxied by their participation in the small-business tax regime *Simples Nacional*),⁵¹ five-digit industry dummies, and municipality dummies.

⁴⁹These 12 amenity proxies, together with indicators for employer location and industry, explain up to 47 percent of the model-implied amenity utility estimates by gender—see Table 7 in the paper.

⁵⁰We define task intensity as the mean z-score of a given task measures across occupations of workers at a given establishment. We obtain task measures for the Brazilian *Classificação Brasileira de Ocupações (CBO)* occupation codes by hand-matching them to US Census occupation codes, which are then linked to the Occupational Information Network (O*NET) task scales constructed by Autor and Dorn (2009) and Acemoglu and Autor (2011).

⁵¹Eligibility for the *Simples Nacional* tax regime requires that the enterprise is a micro- or small business with annual revenues below BRL 1,200,00 (around USD 200,000), that it has no other companies as stakeholders, that it is not internationally owned, that it has no shareholder or partner with significant financial stakes in other companies, and that the enterprise itself has no stake in other companies.

Table E.2. Within-employer gender differentials in amenity utilization

	(1)	(2)	(3)	(4)
	Baseline:	Baseline:	Diff.	Std. err.
	Men	Women	($\times 100$)	($\times 100$)
Indicator: employer provides food stamps	0.513	0.391	0.121***	(0.001)
Share with part-time contract	0.082	0.180	3.426***	(0.003)
Share with hours change since previous year	0.032	0.052	0.481***	(0.001)
Share with paid sick leave	0.039	0.053	1.587***	(0.001)
Share with parental leave	0.000	0.031	3.097***	(0.000)
Share with unpaid leave	0.003	0.004	0.023***	(0.000)
Share with earnings cut since previous year	0.143	0.122	-1.303***	(0.001)
Share with noncontractual-earnings fluctuations	0.969	0.957	-0.318***	(0.001)
Share with work-related accident	0.563	0.471	-1.479***	(0.003)
Share with commute-related accident	0.008	0.005	-0.461***	(0.001)
Share of separations due to firing for unjust reasons	0.005	0.003	-0.110***	(0.000)
Share of separations due to worker death	0.001	0.000	-0.016***	(0.000)
Number of unique establishments	95,948	95,948	95,948	
Observations	9,681,841	7,585,469	17,267,310	

Note: Table shows baseline mean values of 12 amenity proxies by gender (columns [1]–[2]) and estimation results from separate regressions of gender-employer-specific amenity proxies on gender dummies while controlling for employer dummies (columns [3]–[4]). Sample consists of all dual-gender firms in the connected set subject to the same selection criteria as in the main text. Coefficients (column [3]) and standard errors (column [4]) are multiplied by a factor of 100. ***, **, and * denote significance at the 1 percent, 5 percent, and 10 percent levels.

E.6 Robustness for Covariates of Wedge Estimates

We here demonstrate robustness of the analysis of gender wedge covariates from Section 6.6 by projecting estimated gender wedge on a set of covariates as in the main body of the paper but with the variables “female employment share” and “indicator: highest-paid worker is a woman” omitted. Table E.3 shows the result of this analysis. The estimated coefficients are qualitatively the same as in Table 9 of the the main body of the paper, with the notable exceptions that the coefficients on “share of worker separations due to worker death” and “indicator: no major financial stakeholder” are now of the opposite sign. This leaves unchanged our conclusion that our analysis lends limited support to the estimated gender wedge being systematically related to employer-level comparative advantages across genders.

E.7 Basic Solution Algorithm

We start by feeding to the model the estimated labor market parameters $\{\lambda_M^u, \lambda_F^u, s_M^e, s_F^e, s_M^G, s_F^G, \delta_M, \delta_F\}$ and the firm-level estimates of $\{p, \pi_M, \pi_F, z, c_M^{v,0}, c_F^{v,0}\}$. Then, we rank firms according to “composite productivity” \tilde{p}_g for each gender. This is useful because, as stated in Lemma 4, firms that have higher “composite productivity” will pay higher effective wages.

We must first find the equilibrium level of aggregate vacancies V_g . We invert the equation for the

Table E.3. Regression of estimated gender wedges on empl. char.s, by gender (robustness)

	(1)	(2)	(3)
Routine manual task intensity	-0.580*** (0.001)	-0.349*** (0.001)	-0.350*** (0.001)
Nonroutine manual task intensity	1.145*** (0.001)	0.591*** (0.001)	0.596*** (0.001)
Routine cognitive task intensity	0.038*** (0.001)	0.030*** (0.001)	0.017*** (0.001)
Nonroutine cognitive interpersonal task intensity	-0.989*** (0.001)	-0.297*** (0.001)	-0.282*** (0.001)
Nonroutine cognitive analytical task intensity	0.629*** (0.001)	0.276*** (0.001)	0.257*** (0.001)
Share of worker separations due to worker death	-1.697*** (0.008)	0.262*** (0.007)	0.395*** (0.007)
Share of workers with work-related accidents	15.978*** (0.032)	4.579*** (0.026)	4.098*** (0.026)
Indicator: no major financial stakeholders	-0.265*** (0.001)	-0.139*** (0.001)	-0.137*** (0.001)
Industry FEs		✓	✓
Municipality FEs			✓
Number of unique establishments	96,065	96,065	96,065
Observations	17,287,101	17,287,101	17,287,101
R ²	0.274	0.589	0.626

Note: Table reports estimated coefficients from regression of structurally estimated gender wedges on observable establishment characteristics based on equation (27). ***, **, and * denote significance at the 1 percent, 5 percent, and 10 percent levels, respectively.

offer arrival rate from unemployment in (13) to obtain:

$$V_{a,g} = U_{a,g} \left(\frac{\lambda_g^u}{\chi_{a,g}} \right)^{1/\alpha}.$$

Now consider the firm's first order conditions for vacancies in equation (C.6) and the rewritten first-order condition with respect to flow utility in equation (C.10). We use the same transformations that we have used to prove Lemma 4. Define $h_g(\tilde{p}_g) = F(x_g^*(\tilde{p}_g))$. Thus, $h'_g(\tilde{p}_g) = f_g(x_g^*(\tilde{p}_g)) x'_g(\tilde{p}_g)$; therefore, $f(x_g^*(\tilde{p}_g)) = h'_g(\tilde{p}_g)/x'_g(\tilde{p}_g)$. Also, $v(\tilde{p}_g) = V_g h'_g(\tilde{p}_g)/\gamma(\tilde{p}_g)$, so we can rewrite $h'_g(\tilde{p}_g) = v_g(p_g)\gamma(\tilde{p}_g)/V_g$. We assume $c_g^v(v) = c_g^{v,0}v^2/2$. Thus, we rewrite the first-order conditions as

$$\begin{aligned} h'_g(p_g) &= \frac{T_g(\tilde{p}_g - x_g(\tilde{p}_g))}{V_g c_g^{v,0}(p_g)} \left(\frac{1}{\delta + \lambda_g^G + \lambda_g^e(1 - h_g(\tilde{p}_g))} \right)^2 \gamma_g(\tilde{p}_g) \\ x'_g(\tilde{p}_g) &= 2\lambda_g^e \frac{T_g(\tilde{p}_g - x_g(\tilde{p}_g))^2}{V_g c_g^{v,0}(\tilde{p}_g)} \left(\frac{1}{\delta + \lambda_g^G + \lambda_g^e(1 - h_g(\tilde{p}_g))} \right)^3 \gamma_g(\tilde{p}_g), \end{aligned} \quad (\text{E.1})$$

where $\gamma_g(\tilde{p}_g)$ is the density function of \tilde{p}_g . Then, after ranking all firms, we solve for wages and total vacancies posted by integrating the functions above by gender and firm by firm. For every firm j ,

$$\begin{aligned} h'_{g,j} &= \frac{T_g(\tilde{p}_{g,j} - x_{g,j})}{V_g c_{g,j}^{v,0}} \left(\frac{1}{\delta + \lambda_g^G + \lambda_g^e(1 - h_{g,j})} \right)^2 \gamma_g(\tilde{p}_{g,j}) \\ x'_{g,j} &= \frac{2\lambda_g^e T_g(\tilde{p}_{g,j} - x_{g,j})^2}{V_g c_{g,j}^{v,0}} \left(\frac{1}{\delta + \lambda_g^G + \lambda_g^e(1 - h_{g,j})} \right)^3 \gamma_g(\tilde{p}_{g,j}) \\ h_{g,j+1} &= h_{g,j} + h'_{g,j}(\tilde{p}_{g,j+1} - \tilde{p}_{g,j}) \\ x_{g,j+1} &= x_{g,j} + x'_{g,j}(\tilde{p}_{g,j+1} - \tilde{p}_{g,j}), \end{aligned} \quad (\text{E.2})$$

and we calculate total vacancies obtained in equilibrium as $V_g^* = \sum_j v_{g,j} \gamma_g(\tilde{p}_{g,j})(\tilde{p}_{g,j} - \tilde{p}_{g,j-1})$. We solve the algorithm setting the initial conditions $x_{g,0} = \phi_g$ and $h_{g,0} = 0$, and we loop over a multiplier of all vacancy cost shifters $c_{g,j}^{v,0}$ until we obtain that $V_g^* = V_g$. Our solution algorithm produces gender-specific firm-level flow utility $x_{g,j}$, easily converted to wages $w_{g,j} = x_{g,j} - \pi_{g,j}$, gender-specific firm-level recruiting intensities $v_{g,j}$ and gender-specific firm-level ranks in the offer distribution $F_{g,j}$ that are exactly identical to those observed in the data, except for small rounding errors: the correlation between the data and the model-generated data is larger than 0.99999 for all variables, and by definition it's equal to 1 for ranks.

When we perform counterfactuals that involve equalizing female amenities to male amenities, we assume that $c_{j,g}^\pi(\pi_g) = c_{j,g}^{\pi,0} \pi_g^2 / 2$. We estimate the cost shifter $c_{j,g}^{\pi,0}$ by solving this equation for $c_{j,g}^{\pi,0}$ given our estimates of $\pi_{j,g}$. Thus, when a firm j was previously posting amenities $\pi_{j,F}$ for women in equilibrium, and now this level is equalized to the level of men $\pi_{j,M}$, we readjust composite productivity of firm j as follows:

$$\begin{aligned} \tilde{p}_{j,F}^{baseline} &= \tilde{p} - z - c_{j,F}^\pi \frac{\pi_{j,F}^2}{2} \\ \tilde{p}_{j,F}^{counterfactual} &= \tilde{p} - z - c_{j,M}^\pi \frac{\pi_{j,M}^2}{2} \\ &= \tilde{p}_{j,F}^{baseline} + c_{j,F}^\pi \frac{\pi_{j,F}^2}{2} - c_{j,M}^\pi \frac{\pi_{j,M}^2}{2}. \end{aligned}$$

F Simulation Appendix

F.1 Alternative Numerical Solution Algorithm for Simulated Policies

When we simulate policies, we can no longer rely on the model prediction that firms with higher composite productivity $p + \pi_g - c_g^\pi(\pi_g) - \mathbf{1}[g = F]z$ will post higher flow utility and vacancies. The reason is that under the policies we consider, the effective productivity levels of both men and women matter for wages, amenities, and vacancies to be posted for either gender as firms now maximize total profits across markets. Instead, we solve the following firm profit-maximization problem:

$$\begin{aligned} \max_{w, \pi_M, \pi_F, v_m, v_f} \quad & \left\{ T_m v_m (p - w - c_M^\pi(\pi_M)) \left(\frac{1}{\delta_m + \lambda_m^G + \lambda_m^e (1 - F_m(w + \pi_M))} \right)^2 \right. \\ & + T_f v_f (p - z - w - c_F^\pi(\pi_F)) \left(\frac{1}{\delta_f + \lambda_f^G + \lambda_f^e (1 - F_f(w + \pi_F))} \right)^2 \\ & \left. - c_m^v(v_m) - c_f^v(v_f) \right\}, \end{aligned}$$

where the definitions of T_g , F_g and V_g are as in the standard solution algorithm. The only unknowns in this problem are F_M and F_F , two endogenous objects to be determined in the equal-pay policy equilibrium. Firms can still hire both genders, only one gender, or neither.

Denote by \mathcal{F}_g the mapping from $\{F_M, F_F\}$ to the offer distributions implied by firms' behavior. We solve the following system of functional equations:

$$\begin{aligned} \mathcal{F}_M(F_M^*, F_F^*) &= F_M^* \\ \mathcal{F}_F(F_M^*, F_F^*) &= F_F^*, \end{aligned}$$

where $\mathcal{F}_g(F_M^*, F_F^*)$ represents the offer distributions implied by the optimal choices of firms that are a function of the offer distributions in the economy. It's worth noticing that the offer distributions of both genders implicitly depend on the offer distributions of both men and women, because when firms decide which wage to set, they have to take into account the effects this will have for attracting both genders with respect to the competition they face in the ladder.

Therefore, we solve for the equilibrium offer distributions F_M and F_F as follows:

1. Start with a guess for F_M and F_F ; compute the firm's policy functions for optimal wages, amenities, and vacancies, taking F_M and F_F as given.
2. Using equation (11), aggregate optimal vacancies of firms to calculate V_g .
3. Compute the offer distributions F_M and F_F that are implied by the firms' policy functions.
4. Find F_M and F_F such that the offer distributions taken as given by firms and the offer distributions implied by the firms' behavior are identical.

F.2 Counterfactual Simulations under Exogenous Amenities

We perform robustness checks in which we consider amenities as exogenous, so that firms are born with them and post them at zero cost (that is, $c_{j,g}^\pi(\pi_{j,g}) = 0$ for all j, g). Our results remain substantially unchanged for all counterfactuals, because we find that the cost of posting amenities is relatively small for most firms. Results for counterfactuals under the exogenous amenities assumption can be found in Table F.1 below. Results for the counterfactual equal-pay policy under the exogenous amenities assumptions can be found in Table F.2 below.

Table F.1. Equilibrium decomposition of gender pay gap under exogenous amenities

	Baseline	Counterfactuals			
	(0)	(1)	(2)	(3)	(4)
Gender differences in...					
...amenities	✓		✓	✓	
...wedges	✓	✓		✓	
...recruiting costs	✓	✓	✓		
Gender pay gap...	0.074	0.060	0.023	0.018	0.000
...between employers	0.055	0.056	0.046	0.016	0.000
...within employers	0.018	0.004	-0.023	0.002	0.000
Output	1.000	1.001	1.010	1.029	1.034
Worker welfare...	0.000	0.005	0.015	-0.004	0.027
...from total payroll...	0.000	0.010	0.024	0.019	0.029
...for men	0.000	0.000	0.000	0.000	0.000
...for women	0.000	0.010	0.024	0.019	0.029
...from total amenity value...	0.000	-0.005	-0.009	-0.022	-0.002
...for men	0.000	0.000	0.000	0.000	0.000
...for women	0.000	-0.005	-0.009	-0.022	-0.002
Payroll-equivalent welfare change	0.000	0.006	0.018	-0.004	0.033
Employer welfare...	1.000	0.995	1.010	0.986	1.039
...from profits	1.005	1.003	1.010	1.035	1.039
...from wedges	-0.005	-0.008	0.000	-0.049	0.000
Total employment for men	0.757	0.757	0.757	0.757	0.757
Total employment for women	0.760	0.760	0.762	0.760	0.757

Note: Table reports simulation results from model-based counterfactuals with exogenous amenities. Baseline results (column [0]) are compared against counterfactuals without gender differences in amenities (column [1]), gender wedges (column [2]), vacancy posting costs (column [3]), and any gender differences (column [4]).

Table F.2. The equal-pay policy under exogenous amenities

	Baseline	Equal-pay policy
	(0)	(1)
Mean pay for men	1.000	1.000
Mean pay for women	1.000	1.016
Gender pay gap...	0.074	0.057
...between employers	0.055	0.057
...within employers	0.018	0.000
Output	1.000	1.000
Worker welfare...	0.000	0.002
...from total payroll...	0.000	0.002
...for men	0.000	-0.001
...for women	0.000	0.003
...from total amenity value...	0.000	-0.000
...for men	0.000	0.000
...for women	0.000	-0.000
Payroll-equivalent welfare change	0.000	0.003
Employer welfare...	1.000	0.999
...from profits	1.005	1.004
...from wedges	-0.005	-0.005
Total employment for men	0.757	0.757
Total employment for women	0.760	0.759

Note: Table reports results from simulating a counterfactual equal-pay policy when amenities are exogenous. Baseline results (column [0]) are compared against the economy under an equal-pay policy (column [1]).